arXiv:1609.02541v2 [stat.CO] 31 Dec 2016

IMPORTANCE SAMPLING TYPE CORRECTION OF MARKOV CHAIN MONTE CARLO AND EXACT APPROXIMATIONS

MATTI VIHOLA, JOUNI HELSKE, AND JORDAN FRANKS

Abstract. We use an importance sampling (IS) type correction of approximate Markov chain Monte Carlo (MCMC) output to provide consistent estimators. The IS approach, based on unbiased estimators, provides a natural alternative to delayed acceptance (DA) MCMC, but contrary to DA-MCMC, benefits from straightforward parallelisation. We focus on the case where the MCMC simulation is conducted on an approximate marginal distribution, and the IS type correction leads to consistent inference on the corresponding joint distribution. Such a case allows for importance sampling analogues of pseudo-marginal type MCMC methods, such as the grouped independence Metropolis-Hastings and the particle marginal Metropolis-Hastings. We detail strong consistency of the suggested estimators under mild assumptions, and provide central limit theorems with expressions for asymptotic variances. The methods are illustrated in practice with state space models, where our experimental results are promising and show that IS type approach can be favourable against a computationally comparable DA scheme.

1. Introduction

Markov chain Monte Carlo (MCMC) has become a standard tool in Bayesian analysis. The greatest benefit of MCMC is its generality — it is guaranteed to be consistent with virtually no assumptions on the underlying model. However, the practical applicability of MCMC generally depends on the dimension of the unknown variables, the number of data, and the computational resources available. Furthermore, because MCMC is sequential in nature, it can be difficult to implement efficiently with modern parallel and distributed computing architectures; see [38, 56, 94] for general discussion about MCMC in challenging scenarios, and with parallel computing architectures.

In recent years, a number of generic approximation methods have been developed for complex Bayesian inference, for instance variational Bayes [cf. reviews 9, 49], expectation propagation [68] and Laplace approximations [78, 86]. These methods have been reported to provide accurate enough inference for practical purposes in many scenarios, with a fraction of computing time compared with (asymptotically) exact approaches such as MCMC. One problem with approximate methods is that the related bias is typically hard or impossible to quantify, and in practice the only way to ensure the validity of the method is to compare the outcome with an exact method, such as MCMC.

2010 *Mathematics Subject Classification.* Primary 65C60; secondary 60J22, 65C05, 65C35, 65C40.

Key words and phrases. Delayed acceptance, exact approximation, importance sampling, Markov chain Monte Carlo, parallel computing, particle filter, state space model, unbiased estimator.

We promote, in broad terms, a well-known principle: the combination of an approximate method for 'preliminary' inference, and importance sampling type correction of the output. We focus on MCMC targeting an approximate marginal distribution, which is typically based on a marginal likelihood approximation, and consequent importance sampling type estimators for inference on the full posterior. The method is naturally applicable in latent variable models, where MCMC provides a non-parametric approximation of the (non-standard) marginal distribution of the parameter, and the IS-type correction is applied on the latent variable part, which may be closely approximated. The philosophy is therefore similar to the popular integrated nested Laplace approximation [78], but the aim in this paper is exact inference, with the expense of additional computational cost.

Our setting highlights explicitly the connection of importance sampling type correction and recently developed pseudo-marginal type MCMC [3, 57], such as particle MCMC [1], grouped independence Metropolis-Hastings (GIMH) [10], approximate Bayesian computation MCMC [63], the algorithm for estimation of discretely observed diffusions suggested in [11], and annealed importance sampling [50, 69]. It also enables the use of debiased estimators as suggested in [66, 75]. All these methods are often computationally expensive, but our approach is based on independent realisations of such importance sampling type estimates, which can be calculated efficiently in parallel, possibly using distributed resources.

We detail the application of the method with Bayesian state space models, where importance sampling and particle filters are used for correction. We believe that the setup can be useful much more generally, for instance in other types of latent variable models or discretisations of continuous models, such as diffusions or inverse problems.

1.1. **Related work.** Importance sampling correction of MCMC has been suggested early in the MCMC literature; see at least [21, 34, 40]. It has been used, for instance, when estimating Bayes factors using a single MCMC output [22]. Bhattacharya [12] has suggested a consistent estimator based on regeneration and Tan, Doss and Hobert [84] give such estimators in case of using multiple Markov chains.

Using unbiased estimators of importance weights in this context has been suggested at least in [57, 62], who consider a generalisation of the pseudo-marginal method, allowing for likelihood estimators that occasionally take negative values. This allows, for instance, using the recently discovered 'debiasing tricks' [66, 75] in order to construct an unbiased estimator. Similar tricks, which would lead into non-negative unbiased estimators, as required by the pseudo-marginal method, may be unavailable [cf. 45]. Quiroz, Villani and Kohn [73] applied the IS correction in the data sub-sampling context.

Nested or compound sampling has also appeared in many forms in the Monte Carlo literature. The SMC^2 algorithm [14] is based on an application of nested sequential Monte Carlo steps, which has similarities with our framework, and the IS^2 method [87] focuses on the case where the preliminary inference is based on importance sampling. We focus on the MCMC approximation of the marginal distribution, which we believe often to be easily implementable in practice. The

Markov chain nature of the marginal Monte Carlo approximation comes also with some extra theoretical issues which we address.

Another specific feature of our work is that we connect IS type correction explicitly with recently discovered 'exact approximation MCMC' techniques [1, 3, 10] and delayed acceptance [16]. Theoretical advances [2, 5, 6, 15, 27, 58, 82] have already led to more efficient implementation of such methods, but has also revealed fundamental limitations. For instance, the methods may suffer from slow (non-geometric) convergence in practically interesting scenarios [3, 54]. Adding dependence to the estimators [cf. 6], such as using the recently proposed correlated version of the pseudo-marginal MCMC [19], may help in more efficient implementation. However, a successful implementation of such a method may not always be possible, and the question of efficient parallelisability remains a challenge. The blocked parallelisable particle Gibbs [83] has appealing limiting properties, but its implementation still requires synchronisation between every update cycle, which may be costly in some computing environments. We recognise that our IS approach comes with its own limitations (see Section 10), but we believe that it can often be very useful, especially with parallel computing facilities.

1.2. **Outline.** After preliminaries and notation in Section 2, we discuss a general importance sampling (IS) type correction of MCMC outputs and we state related consistency results in Section 3. We detail the general case where unbiased estimators are used (Proposition 1) and then focus on the setting in order to allow inference in an extended state space (Theorem 4). The latter result is based on a general concept (Assumption 3), which we call 'proper weighting scheme' following the terminology of Liu [61]. It accommodates naturally estimators from importance sampling and particle filters. We believe that Assumption 3, or the more general Assumption 2, applies in a wide variety of scenarios.

Section 4 discusses conditions under which our estimators admit a central limit theorem. The expressions for the asymptotic variance may be useful, for instance, when optimising the computational resources. Section 5 focuses on estimators that use the so-called jump chain representation [20, 24, 27], and apply IS correction for each accepted state. In this context, it is possible to employ variance reduction techniques, such as stratification or Rao-Blackwellisation using the scheme of [24]. Section 6 starts by making an explicit connection between the proper weighting schemes and the pseudo-marginal type MCMC, and then between the IS type correction of MCMC and a two-stage delayed acceptance scheme [16].

In Section 7, we discuss how our importance sampling correction method applies in general state space models (SSM), and detail how importance sampling and particle filters lead to proper weighting schemes. Section 8 discusses an important sub-class of state space models, where the state dynamics are linear-Gaussian. We detail how a Laplace approximation can be used in this context, providing approximate preliminary inference, and efficient importance sampling and particle filtering schemes. We compare several algorithmic variations in Section 9 with Poisson observations and with a stochastic volatility model. We conclude in Section 10 with a discussion of our findings.

2. Notation and preliminaries

Throughout the paper, we consider general state spaces while using standard integral notation. If the model at hand is defined on Euclidean space using standard probability densities, the following paragraph can be skipped.

Each space X is assumed to be equipped with a σ-finite dominating measure '$\mathrm{d}x$' on a σ-algebra denoted with a corresponding calligraphic letter, such as $\mathcal{X}$. Product spaces are equipped with the related product σ-algebras and product dominating measures. If X is a subset of an Euclidean space $\mathbb{R}^d$, $\mathrm{d}x$ is taken by default as the Lebesgue measure and $\mathcal{X}$ as the Borel subsets of X.

If μ is a probability density on X, we denote the support of μ as $\mathrm{supp}(\mu) := \{x \in \mathsf{X} : \mu(x) > 0\}$, and the probability measure corresponding to μ with the same symbol $\mu(\mathrm{d}x) := \mu(x)\mathrm{d}x$. If $g : \mathsf{X} \to \mathbb{R}$, we denote $\mu(g) := \int g(x)\mu(\mathrm{d}x)$, whenever well-defined. For a probability density or measure μ on X and $p \in [1, \infty)$, we denote by $L^p(\mu)$ the set of measurable $g : \mathsf{X} \to \mathbb{R}$ with $\mu(|g|^p) < \infty$, and by $L^p_0(\mu) := \{g \in L^p(\mu) : \mu(g) = 0\}$ the corresponding set of zero-mean functions. If P is a Markov transition probability, we denote the probability measure $(\mu P)(A) := \int \mu(\mathrm{d}x)P(x, A)$, and the function $(Pg)(x) := \int P(x, \mathrm{d}y)g(y)$. Iterates of transition probabilities are defined recursively through $P^n(x, A) := \int P(x, \mathrm{d}y)P^{n-1}(y, A)$ for $n \geq 1$, where $P^0(y, A) := \mathbf{1}(y \in A)$.

We follow the convention $0/0 := 0$, and if $a, b \in \mathbb{N}$, we denote by $a{:}b$ the integers within the interval $[a, b]$. We use this notation in indexing, so that $x_{a:b} = (x_a, \ldots, x_b)$, $x^{(a:b)} = (x^{(a)}, \ldots, x^{(b)})$. If $a > b$, then $x_{a:b}$ or $x^{(a:b)}$ is void, so that for example $g(x, y_{1:0})$ means $g(x)$. Similarly, if $i_{1:T}$ is a vector, then $x^{(i_{1:T})} = (x^{(i_1)}, \ldots, x^{((i_T)})$ and $x_{1:T}^{(i_{1:T})} = (x_1^{(i_1)}, \ldots, x_T^{(i_T)})$. We also use double-indexing, such as $x_k^{(1:m,1:n)} = (x_k^{(1,1)}, \ldots, x_k^{(1,m)}, x_k^{(2,1)}, \ldots, x_k^{(m,n)})$.

Throughout the paper, we assume the underlying MCMC scheme to be Harris ergodic, as below. Recall that virtually all MCMC schemes are Harris [cf. 70, 85], although in some cases careless implementation could lead to a non-Harris chain [cf. 76].

Definition 1 (Harris ergodicity)**.** The Markov chain $(X_k)_{k\geq 1}$ on X is Harris recurrent with invariant probability μ, if for every $g \in L^1(\mu)$, and every initial distribution, the ergodic averages are strongly consistent,

$$\frac{1}{n}\sum_{k=1}^{n} g(X_k) \xrightarrow{n\to\infty} \mu(g), \qquad \text{almost surely.}$$

We call such a chain *Harris ergodic (with respect to μ).*

3. General importance sampling type correction of MCMC

Hereafter, π_a is a probability density on X and represents an approximation of a probability density π of interest. The consistency of IS-type correction relies on the following mild assumption.

Assumption 1. The density π_a and the related Markov chain $(X_k)_{k\geq 1}$ satisfy

(i) $(X_k)_{k\geq 1}$ is Harris ergodic with respect to π_a.
(ii) $\mathrm{supp}(\pi) \subset \mathrm{supp}(\pi_a)$.
(iii) $w_u(x) := c_w\pi(x)/\pi_a(x)$, where $c_w > 0$ is a constant.

If Assumption 1 holds and it is possible to calculate the unnormalised importance weight $w_u(x)$ pointwise, the chain $(X_k)_{k\geq 1}$ can be weighted in order to approximate $\pi(g)$ for every $g \in L^1(\pi)$, using (self-normalised) importance sampling [e.g. 21, 34]

$$\frac{\sum_{k=1}^n w_u(X_k)g(X_k)}{\sum_{j=1}^n w_u(X_j)} = \frac{n^{-1}\sum_{k=1}^n w_u(X_k)g(X_k)}{n^{-1}\sum_{j=1}^n w_u(X_j)} \xrightarrow[\text{a.s.}]{n\to\infty} \frac{\pi_a(w_u g)}{\pi_a(w_u)} = \pi(g), \tag{1}$$

as the Harris ergodicity guarantees the almost sure convergence of both the numerator and the denominator.

In case of marginal inference which we detail later, the importance weights $w_u(x)$ are defined in terms of marginal densities and the functions $g(x)$ are conditional expectations, which are impossible or extremely costly to evaluate. Instead, it is often possible to construct unbiased estimators, which may be used in place of $w_u(X_k)$ and $g(X_k)$, under mild conditions. In order to formalise such a setting, we formulate the following abstract setting, which accommodates further generalisations of importance sampling type corrections.

Assumption 2. Suppose Assumption 1 holds, let $g \in L^1(\pi)$ and let $S_k = \big(W_k, \zeta_k\big)$ be $\mathbb{R}^2$-valued random variables conditionally independent given $(X_k)_{k\geq 1}$, such that the distribution of S_k depends only on the value of X_k, and satisfies

$$\mathbb{E}[W_k\zeta_k \mid X_k = x] = g(x)w_u(x) \qquad \text{and} \qquad \mathbb{E}[W_k \mid X_k = x] = w_u(x),$$

for π_a-almost every $x \in \mathsf{X}$, and such that

$$\pi_a(m^{(1)}) < \infty, \qquad \text{where} \qquad m^{(1)}(x) := \mathbb{E}\big[|W_k|(1+|\zeta_k|) \;\big|\; X_k = x\big].$$

We record the following simple statement which guarantees consistency under Assumption 2.

Proposition 1. *If Assumption 2 holds, then*

$$E_n(g) := \frac{\sum_{k=1}^n W_k\zeta_k}{\sum_{i=1}^n W_i} \xrightarrow[\text{a.s.}]{n\to\infty} \pi(g).$$

The proof of Proposition 1 follows by observing that $(X_k, S_k)_{k\geq 1}$ is Harris-ergodic and the functions $h_1(x, w, \zeta) = w\zeta$ and $h_2(x, w, \zeta) = w$ are integrable with respect to its invariant distribution

$$\check{\pi}(\mathrm{d}x \times \mathrm{d}s) := \pi_a(\mathrm{d}x)Q(x, \mathrm{d}s), \tag{2}$$

where $Q(x, A) := \mathbb{P}(S_k \in A \mid X_k = x)$; see Lemma 21 in Appendix A.

In the specific cases we consider later, we define $W_k := U_k/\pi_a(X_k)$, where $(U_k, \zeta_k)_{k\geq 1}$ are random variables which satisfy

$$\mathbb{E}[U_k\zeta_k \mid X_k = x] = c_w\pi(x)g(x) \qquad \text{and} \qquad \mathbb{E}[U_k \mid X_k = x] = c_w\pi(x).$$

Such assumptions are also required by pseudo-marginal type MCMC algorithms, with the additional requirement that $W_k \geq 0$; see Section 6. In some cases, $\pi_a(x)$ may be intractable, for instance if (X_k) is itself a pseudo-marginal algorithm targeting π_a. It may still be possible to estimate the ratio $w_u(x)$ in an unbiased fashion.

Before focusing on our main topic on an extended state space, let us consider a case where only the weights $w_u(X_k)$ are replaced with unbiased estimates W_k. This idea has been suggested in more specific contexts earlier at least in [62, 73].

Proposition 2. *Suppose $g \in L^1(\pi)$, Assumption 1 holds, $(W_k)_{k\geq 1}$ are independent random variables on $\mathbb{R}$ conditional on $(X_k)_{k\geq 1}$, the distribution of W_k depends only on the value of X_k, and satisfies $\mathbb{E}[W_k \mid X_k = x] \overset{\text{a.s.}}{=} w_u(x)$ for π_a-almost every $x \in \mathsf{X}$. If*

$$\pi_a\big(m_W^{(1)}(|g|+1)\big) < \infty, \qquad \text{where} \qquad m_W^{(1)}(x) := \mathbb{E}\big[|W_k| \;\big|\; X_k = x\big], \tag{3}$$

then the IS estimator is strongly consistent,

$$\frac{\sum_{k=1}^n g(X_k)W_k}{\sum_{i=1}^n W_i} \xrightarrow[\text{a.s.}]{n\to\infty} \pi(g).$$

In particular, (3) *holds if $W_k \geq 0$ almost surely.*

The case $W_k \geq 0$ coincides with the case where pseudo-marginal algorithms can be implemented, and provide similarly consistent estimators for all $g \in L^1(\pi)$. More generally, if the relative error in W_k is bounded, that is, there exists $c < \infty$ such that $m_W^{(1)}(x) \leq c w_u(x)$ for all $x \in \mathsf{X}$, then (3) holds for all $g \in L^1(\pi)$. In general, however, we believe that the estimator may fail to be consistent for all $g \in L^1(\pi)$, and an extra condition such as (3) might be indeed necessary.

We now focus on a more specific setting, which is of our main interest. The aim is inference of a joint target density $\pi^*(x,z) := \pi(x)r(z \mid x)$ on an extended state space $\mathsf{X} \times \mathsf{Z}$. For every function $f \in L^1(\pi^*)$, we denote by $f^*(x) := \int r(z \mid x) f(x,z) \mathrm{d}z$ the conditional expectation of f given x, and $\pi^*(f) = \pi(f^*)$. The following guarantees Assumption 2 with $g = f^*$, and therefore consistency for a class of functions $f \in \mathcal{L} \subset L^1(\pi^*)$.

Assumption 3 (Proper weighting scheme)**.** Suppose Assumption 1 holds, and

(i) $(W_k, V_k^{(1:m)}, Z_k^{(1:m)})_{k\geq 1}$ are conditionally independent random variables on $\mathbb{R}\times[0,1]^m \times \mathsf{Z}^m$ given $(X_k)_{k\geq 1}$, such that the distributions of $(W_k, V_k^{(1:m)}, Z_k^{(1:m)})$ only depend on X_k.

(ii) $\mathbb{E}[W_k \mid X_k = x] = w_u(x)$ for π_a-almost every $x \in \mathsf{X}$.

(iii) $\mathbb{E}[W_k \xi_k(f) \mid X_k = x] = w_u(x) f^*(x)$ for all $f \in \mathcal{L}$ and π_a-almost every $x \in \mathsf{X}$, where

$$\xi_k(f) := \sum_{i=1}^m V_k^{(i)} f(X_k, Z_k^{(i)}).$$

We call $(W_k, V_k^{(1:m)}, Z_k^{(1:m)})_{k\geq 1}$, or equivalently $\big(W_k, \xi_k(\,\cdot\,)\big)_{k\geq 1}$, a proper weighting scheme.

Remark 3. Regarding Assumption 3:

(i) The number of 'samples' m may also be (a random variable) dependent on $X_k = x$. This extension is trivial, but we chose not to cover it explicitly for the sake of more concise notation.

(ii) In fact, (i) fits into a straightforward generalisation which assumes ξ_k to be *random probability measures* on Z^m such that $\xi_k(f) := \int f(X_k, z)\mathrm{d}z$, requiring only $\mathbb{E}[W_k \xi_k(f) \mid X_k = x] = w_u(x) f^*(x)$ for all $f \in \mathcal{L}$.

(iii) In some cases, such as debiasing [66, 75], even further generalisation which allows ξ_k to be *random signed measures* may be appropriate. We do not, however, consider such a setting in detail here.

Note that Assumption 3 implies that $\mathbb{E}[W_k] = w_u(x)$ by taking $f \equiv 1$. We record the following consistency result, which follows from Proposition 1.

Theorem 4. *Suppose Assumption 3 holds and $f \in \mathcal{L}$. Then,*

$$E_n(f) := \frac{\sum_{k=1}^n W_k \xi_k(f)}{\sum_{j=1}^n W_j}, \tag{4}$$

satisfies $E_n(f) \xrightarrow{n\to\infty} \pi^*(f)$ almost surely, if

$$\pi_a(m^{(1)}) < \infty, \qquad \text{where} \qquad m^{(1)}(x) := \mathbb{E}\big[|W_k|(1+|\xi_k(f)|) \;\big|\; X_k = x\big]. \tag{5}$$

In particular, if $W_k \geq 0$ almost surely for all $k \geq 1$, then (5) holds.

Let us next exemplify how a proper weighting scheme stems from (unnormalised) importance sampling, by noticing that

$$w_u(x) f^*(x) = c_w \frac{\int \pi^*(z,x) f(x,z) \mathrm{d}z}{\pi_a(x)}.$$

Proposition 5. *Suppose Assumption 1 holds and $q^{(x)}(\,\cdot\,)$ defines a probability density on Z for each $x \in \mathsf{X}$ and $\mathrm{supp}(\pi^*) \subset \{(x,z) : \pi_a(x) q^{(x)}(z) > 0\}$. Let $Z_k^{(1:m)} \overset{\text{i.i.d.}}{\sim} q^{(X_k)}$ and denote*

$$\omega_k^{(i)} := \frac{c_w \pi^*(X_k, Z_k^{(i)})}{q^{(X_k)}(Z_k^{(i)})}, \qquad W_k := \frac{\omega_k^*}{m \pi_a(X_k)}, \qquad \omega_k^* := \sum_{i=1}^m \omega_k^{(i)},$$

where $c_w > 0$ is some constant, and let $V_k^{(i)} := \frac{\omega_k^{(i)}}{\omega_k^*} \mathbf{1}\,(\omega_k^* > 0) + \frac{1}{m} \mathbf{1}\,(\omega_k^* = 0)$.

Then, $(W_k, V_k^{(1:m)}, Z_k^{(1:m)})_{k\geq 1}$ forms a proper weighting scheme with $\mathcal{L} = L^1(\pi^*)$.

Because $W_k \geq 0$ in Proposition 5, then the estimator is consistent for all $f \in L^1(\pi^*)$, and simplifies into the form

$$\frac{\sum_{k=1}^n \pi_a^{-1}(X_k) \sum_{i=1}^m \omega_k^{(i)} f(X_k, Z_k^{(i)})}{\sum_{k=1}^n \pi_a^{-1}(X_k) \sum_{i=1}^m \omega_k^{(i)}} \xrightarrow[\text{a.s.}]{n\to\infty} \pi^*(f).$$

We conclude this section by noticing that proper weighting property is inherited in sub-sampling, and in convex combinations when the weights are non-negative. The first is immediate.

Proposition 6. *Suppose that $(W_k, V_k^{(1:m)}, Z_k^{(1:m)})_{k\geq 1}$ forms a proper weighting scheme. Let (I_k) be random variables conditionally independent of $(X_k, Z_k^{(i)})$ such that $\mathbb{P}(I_k = i) = V_k^{(i)}$. Then, $(W_k, 1, Z_k^{(I_k)})_{k\geq 1}$ forms a proper weighting scheme.*

Sub-sampling leads to the following simple estimator, which is consistent under the conditions stated in Theorem 4,

$$\frac{\sum_{k=1}^n W_k f(X_k, Z_k^{(I_k)})}{\sum_{k=1}^n W_k} \xrightarrow[\text{a.s.}]{n\to\infty} \pi^*(f)$$

The extra sub-sampling randomisation leads to a higher asymptotic variance, but may be preferable if the weighted sample $\big(W_k, (X_k, Z_k^{(I_k)})\big)_{k\geq 1}$ needs to be stored for further analysis.

We conclude by recording the following statement about convex combinations.

Proposition 7. *Suppose* $\Pi_{k,j} := (W_{k,j}, \xi_{k,j})$ *and* $(\Pi_{k,j})_{k\geq 1}$ *forms a proper weighting scheme for each* $j \in \{1{:}N\}$*, the weights* $W_{k,j} \geq 0$ *almost surely for all* k, j*, and the conditional joint laws of* $(\Pi_{k,1}, \dots, \Pi_{k,N})$ *depend only on* X_k*. Then,*

$$W_k := \sum_{j=1}^{N} \beta_j W_{k,j} \quad \textit{and} \quad \xi_k := \sum_{j=1}^{N} \beta_j \tilde{W}_{k,j} \xi_{k,j},$$

where $\tilde{W}_{k,j} = W_{k,j}/W_k$ *if* $W_k > 0$ *and, say,* $\tilde{W}_{k,j} = 1/N$*, otherwise, define a proper weighting scheme.*

4. Asymptotic variance and a central limit theorem

The asymptotic variance is a common efficiency criterion for Markov chains. When a central limit theorem (CLT) holds, the limiting variance coincides with the asymptotic variance, which also appears in estimators discussed in Section 3.

Definition 2. Suppose the Markov chain $(X_k)_{k\geq 1}$ on X has transition probability P which is Harris ergodic with respect to invariant probability π_a. For $f \in L^2(\pi_a)$, the asymptotic variance of f with respect to P is

$$\mathrm{Var}(f, P) := \lim_{n\to\infty} \mathbb{E}\left(\frac{1}{\sqrt{n}} \sum_{k=1}^{n} \big[f(X_k^{(s)}) - \pi_a(f) \big] \right)^2,$$

whenever the limit exists in $[0, \infty]$, where $(X_k^{(s)})_{k\geq 1}$ stands for the *stationary Markov chain* with transition probability P, that is, with $X_1^{(s)} \sim \pi_a$.

In what follows, we denote by $\bar{f}(x, z) = f(x, z) - \pi^*(f)$ the centred version of any $f \in L^1(\pi^*)$.

Theorem 8. *Suppose that the conditions of Theorem 4 are satisfied, let* $f \in \mathcal{L} \cap L^2(\pi^*)$*, denote* $\bar{f}(x, z) := f(x, z) - \pi^*(f)$ *and assume*

$$\pi_a(m^{(2)}) < \infty, \qquad \textit{where} \qquad m^{(2)}(x) := \mathbb{E}\big[|W_k \xi_k(\bar{f})|^2 \;\big|\; X_k = x \big].$$

If either of the following hold:

(i) $(X_k)_{k\geq 1}$ *is reversible and* $\mathrm{Var}(w_u \bar{f}^*, P) < \infty$*, or*

(ii) $\sum_{n=1}^{\infty} n^{-3/2} \big\{ \pi\big(\big[\sum_{k=0}^{n-1} P^k w_u \bar{f}^* \big]^2 \big) \big\}^{1/2} < \infty$*,*

then, the estimator $E_n(f)$ *defined in* (4) *satisfies a CLT*

$$\sqrt{n}[E_n(f) - \pi^*(f)] \xrightarrow{n\to\infty} N\left(0, \frac{\mathrm{Var}(w_u \bar{f}^*, P) + \pi_a(v)}{c_w^2} \right)$$

in distribution, where $v(x) := \mathrm{Var}\big(W_k \xi_k(\bar{f}) \;\big|\; X_k = x \big)$.

Proof of Theorem 8 is given in Appendix B.

In case of reversible chains, the condition in Theorem 8 (i) is essentially optimal, and the CLT relies on a result due to Kipnis and Varadhan [51]. The condition always holds when $(X_k)_{k\geq 1}$ is geometrically ergodic, for instance $(X_k)_{k\geq 1}$ is a random-walk Metropolis algorithm and π_a is light-tailed [46, 77]. In case $(X_k)_{k\geq 1}$ is sub-geometric, such as polynomial, extra conditions are required; see for instance [47]. The condition (ii) which applies for non-reversible chains is also nearly optimal, and relies on a result due to Maxwell and Woodroofe [64]. See also the review on Markov chain CLTs by Jones [48].

Note that the latter term $\pi_a(v)$ in the asymptotic variance expression contains the contribution of the 'noise' in the IS estimates. If the estimators $W_k\xi_k(f)$ are made increasingly accurate, in the sense that $\pi_a(v)$ becomes negligible, the limiting case corresponds to an IS corrected approximate MCMC, and evaluating the conditional expectation $f^*(x)$. We conclude this section by noticing that the contribution of IS noise can be estimated in a straightforward manner.

Theorem 9. *Suppose conditions of Theorem 4 are satisfied and $f \in \mathcal{L} \cap L^2(\pi^*)$. Suppose also that $\pi_a(v) < \infty$ where v is defined in Theorem 8, and $\pi_a(m_W^{(2)}) < \infty$, where $m_W^{(2)}(x) := \mathbb{E}[W_k^2 \mid X_k = x]$. Then, the estimator*

$$v_n := \sum_{k=1}^{n} \big(\tilde{W}_k^{(n)}\big)^2 \big(\xi_k(f) - E_n(f)\big)^2 \quad \textit{where} \quad \tilde{W}_k^{(n)} := \frac{W_k}{\sum_{j=1}^n W_j}$$

satisfies $nv_n \to \pi_a(v)/c_w^2$ almost surely as $n \to \infty$.

Proof of Theorem 9 is given in Appendix B.

The estimator v_n in Theorem 9 provides a lower bound for the CLT variance σ_f^2/n of $E_n(f)$. It can provide useful information about the importance sampling noise contribution, and may be used as an optimisation criteria when adjusting the accuracy of the related estimators. For instance, it may be compared with a Markov chain asymptotic variance estimator (see, e.g., the review [31] and references therein), which includes the marginal chain contribution $\mathrm{Var}(w_u\bar{f}^*, P)/c_w^2$.

5. Jump chain estimators

Many MCMC algorithms include an accept-reject mechanism, which results in blocks of repeated values $X_k = \ldots = X_{k+b}$. When $w_u(X_k)f^*(X_k)$ are estimated, an obvious approach is to construct only one estimator for each accepted state. This may allow for variance reduction, for instance when replacing simple independent importance sampling with block-wide stratified estimator, or using particle filter estimators as in Proposition 17, and choosing the number of particles proportional to the block length.

To formalise such an algorithm we consider the "jump chain" representation of $(X_k)_{k\geq 1}$ [cf. 24].

Definition 3 (Jump chain)**.** Suppose that $(X_k)_{k\geq 1}$ is a Markov chain with non-degenerate transition probability P, that is, $P(x, \mathsf{X} \setminus \{x\}) > 0$ for all $x \in \mathsf{X}$. The jump chain $(\tilde{X}_k)_{k\geq 1}$ with corresponding durations $(N_k)_{k\geq 1}$ of $(X_k)_{k\geq 1}$ is defined as formally as follows: $\tilde{X}_1 \equiv X_1$ and $N_1 := \inf\{k \geq 1 : X_{k+1} \neq \tilde{X}_1\}$, and then recursively

$$\tilde{X}_i := X_{\bar{N}_{i-1}+1} \quad \text{and} \quad N_i := \inf\big\{k \geq 1 : X_{\bar{N}_{i-1}+k+1} \neq \tilde{X}_{i-1}\big\},$$

where $\bar{N}_i := \sum_{j=1}^{i} N_j$.

Remark 10. We note the following observations regarding Definition 3:

(i) Definition 3 is consistent for Harris ergodic chains P, because $P(x, \{x\}) < 1$ for all $x \in \mathsf{X}$, unless π_a is trivial, that is, $\pi_a(\{x\}) = 1$ for some $x \in \mathsf{X}$.

(ii) If $(X_k)_{k\geq 1}$ corresponds to a Metropolis-Hastings chain, with non-diagonal proposal distributions q (that is, $q(x,\{x\}) = 0$ for every $x \in \mathsf{X}$), then the jump chain $(\tilde{X}_k)$ consists of the accepted states, and $N_k - 1$ is the number of rejections occurred at state $(\tilde{X}_k)$.

Hereafter, we denote by $\alpha(x) := \mathbb{P}(X_{k+1} \neq X_k \mid X_k = x)$ the overall acceptance probability at x. The jump chain estimators which we consider are of the following general form:

Assumption 4. Suppose that Assumption 1 holds, let $(\tilde{X}_k)_{k\geq 1}$ denote the corresponding jump chain (Definition 3), and assume $f^* \in L^1(\pi)$. Suppose that $S_k = (\tilde{W}_k, \tilde{\zeta}_k)$ are conditionally independent random variables on $\mathbb{R}^2$ given $(\tilde{X}_k)_{k\geq 1}$, whose distribution depends only on the value of $\tilde{X}_k$, and which satisfy

$$\mathbb{E}[\tilde{W}_k\tilde{\zeta}_k \mid \tilde{X}_k = x] = \frac{w_u(x)f^*(x)}{\alpha(x)} \qquad \text{and} \qquad \mathbb{E}[\tilde{W}_k \mid \tilde{X}_k = x] = \frac{w_u(x)}{\alpha(x)},$$

for π_a-almost every $x \in \mathsf{X}$, and

$$\pi_a(\alpha\tilde{m}^{(1)}) < \infty, \qquad \text{where} \qquad \tilde{m}^{(1)}(x) := \mathbb{E}\big[|\tilde{W}_k|(1 + |\tilde{\zeta}_k|) \;\big|\; \tilde{X}_k = x\big].$$

Proposition 11. *Suppose Assumption 4 holds, then*

$$E_n(f^*) := \frac{\sum_{k=1}^n \tilde{W}_k\tilde{\zeta}_k}{\sum_{i=1}^n \tilde{W}_k} \xrightarrow[\text{a.s.}]{n\to\infty} \pi(f^*).$$

The proof follows from Proposition 1, because $(\tilde{X}_k)$ is Harris ergodic with invariant probability $\tilde{\pi}(\mathrm{d}x) = \pi_a(x)\alpha(x)/\pi_a(\alpha)$; see Proposition 24 in Appendix C.

We consider next an intuitive jump chain estimator based on the durations $(N_k)_{k\geq 1}$ and a proper weighting scheme.

Assumption 5. Suppose that Assumption 1 holds, and let $(\tilde{X}_k, N_k)_{k\geq 1}$ denote the corresponding jump chain (Definition 3). Let $(W_k, \xi_k)_{k\geq 1}$ be a proper weighting scheme (Assumption 3), but such that the variables in the scheme may depend on both $\tilde{X}_k$ and N_k, and satisfy

$$\mathbb{E}[W_k\xi_k(f) \mid \tilde{X}_k = x, N_k = n] = w_u(x)f^*(x),$$

for all $f \in \mathcal{L}$ and π_a-almost every $x \in \mathsf{X}$.

Theorem 12. *Suppose Assumption 5 holds, and*

$$\pi_a(b) < \infty, \qquad \textit{where} \qquad b(x) := \sup_{n\geq 1} \mathbb{E}\big[|W_k|(1+|\xi_k(f)|) \;\big|\; \tilde{X}_k = x, N_k = n\big]. \tag{6}$$

Then,

$$E_n(f) := \frac{\sum_{k=1}^n N_k W_k \xi_k(f)}{\sum_{i=1}^n N_k W_k} \xrightarrow[\text{a.s.}]{n\to\infty} \pi^*(f). \tag{7}$$

The proof follows from Proposition 11, because the holding times $N_k \geq 0$ are conditionally independent geometric random variables with parameter $\alpha(\tilde{X}_k)$; see Proposition 24 (ii) in Appendix C. Therefore $\mathbb{E}[N_k \mid \tilde{X}_k = x] = 1/\alpha(x)$, and so $\tilde{\zeta}_k = \xi_k(f)$ and $\tilde{W}_k = N_k W_k$ satisfy Assumption 4.

Remark 13. Regarding Assumption 5:

(i) The condition (6) is practically convenient, but not optimal. For instance, it is easy to find examples where the estimator is strongly consistent, even though $b(x) = \infty$. In practice, the estimators are usually chosen either as independent of n, or increasingly accurate in n; see the discussion below.
(ii) In case $\xi_k(f)$ are constructed as a convex combination of N_k independent estimators, the jump chain estimator coincides with the simple estimator discussed in Section 4. However, the jump chain estimator offers more flexibility which may allow variance reduction, for instance by using stratification.
(iii) Even though we believe that the estimators of the form (7) are often appropriate, we note that general form of the estimator as given in Assumption 4 may be useful in some cases, in order to further reduce variance. For instance, Assumption 4 accommodates the case where Rao-Blackwellised lower-variance estimators of $1/\alpha(\tilde{X}_j)$ are used instead of N_j, as suggested in [24].

Let us finally consider the central limit theorem corresponding the estimator in Theorem 12.

Theorem 14. *Suppose Assumption 5 and condition* (6) *hold, let* $f \in \mathcal{L} \cap L^2(\pi^*)$, *and assume*

$$\pi_a\big(\alpha \tilde{m}^{(2)}\big) < \infty, \qquad \text{where} \qquad \tilde{m}^{(2)}(x) := \mathbb{E}\big[N_k^2 W_k^2 \xi_k^2(\bar{f}) \;\big|\; \tilde{X}_k = x\big], \tag{8}$$

and one of the following:

(i) $(X_k)_{k\geq 1}$ *is reversible and* $\mathrm{Var}(w_u \bar{f}^*, P) < \infty$.
(ii) There exists $g \in L^2(\pi_a)$ *which satisfies the Poisson equation* $g - Pg = w_u \bar{f}^*$.

Then, the estimator $E_n(f)$ *in* (7) *satisfies*

$$\sqrt{n}\big[E_n(f) - \pi^*(f)\big] \xrightarrow{n\to\infty} N(0, \sigma^2) \qquad \text{in distribution,}$$

where the limiting variance can be given as:

$$\sigma^2 = \frac{\pi_a(\alpha)}{c_w^2}\Big[\mathrm{Var}\big(w_u \bar{f}^*, P\big) + \pi_a(\alpha \tilde{v})\Big], \tag{9}$$

where $\tilde{v}(x) := \mathbb{E}\big[N_k^2 \mathrm{Var}\big(W_k \xi_k(\bar{f}) \;\big|\; \tilde{X}_k = x, N_k\big) \;\big|\; \tilde{X}_k = x\big]$.

Proof of Theorem 14 is given in Appendix C.

Let us briefly discuss the conditions and implications of Theorem 14 under certain specific cases.

(i) In the common case where the acceptance probability is bounded from below, $\inf_x \alpha(x) > 0$, the scheme W_k and ξ_k may be safely taken to be independent of N_k. Namely, (8) may be replaced with

$$\pi_a(b) < \infty, \qquad \text{where} \qquad b(x) = \sup_{n\geq 1} \mathbb{E}\big[W_k^2 \xi_k^2(\bar{f}) \;\big|\; \tilde{X}_k = x, N_k = n\big],$$

because $\tilde{v}(x) \leq \tilde{m}^{(2)}(x) \leq \frac{2-\alpha(x)}{\alpha^2(x)} b(x)$.
(ii) When W_k and ξ_k correspond to $m \propto N_k$ independent estimates, we typically have

$$\mathrm{Var}\big(W_k^2 \xi_k^2(\bar{f}) \;\big|\; \tilde{X}_k = x, N_k\big) \leq \frac{\hat{v}(x)}{N_k}.$$

Then, $\pi_a(\alpha\tilde{v}) \leq \pi_a(\hat{v})$, which leads to a similar asymptotic variance as with simple IS correction (Theorem 8).

(iii) If $(X_k)_{k\geq 1}$ is L^2-geometrically ergodic, then the acceptance probability is (essentially) bounded away from zero [77], and $g := \sum_{k\geq 0} P^k w_u \bar{f}^* \in L^2(\pi_a)$ satisfies $g - Pg = w_u \bar{f}^*$.

Remark 15. Our CLT applies in the non-reversible only when a solution $g \in L^2(\pi_a)$ to the Poisson equation $g - Pg = w_u \bar{f}^*$ exists. We believe that the result holds more generally, but this would requires showing that the jump chain $(\tilde{X}_k)_{k\geq 1}$ inherits a central limit theorem from the base chain $(X_k)_{k\geq 1}$ under more general conditions.

6. Connection with delayed acceptance pseudo-marginal MCMC

We start by making an explicit connection between IS correction of MCMC and pseudo-marginal type MCMC algorithms.

Assumption 6. Suppose that $f^* : \mathsf{X} \to \mathbb{R}$ and $Q(x, \cdot)$ defines a probability distribution on $\mathbb{R}_+ \times \mathbb{R}$ for each $x \in \mathsf{X}$, such that $(U, \zeta) \sim Q(x, \cdot)$ satisfy

$$\mathbb{E}[U] = \varpi_u(x), \qquad \mathbb{E}[U\zeta] = \varpi_u(x) f^*(x).$$

Algorithm 1. Suppose Assumption 6 holds with $\varpi \propto \pi$ and $q(x, \cdot)$ defines a probability density on X for every $x \in \mathsf{X}$. Let $X_0 \in \operatorname{supp}(\pi)$, $U_0 > 0$, $\zeta_0 \in \mathbb{R}$, and for $k \geq 1$ iterate:

(i) Draw $\tilde{X}_k \sim q(X_{k-1}, \cdot)$ and $(\tilde{U}_k, \tilde{\zeta}_k) \sim Q(\tilde{X}_k, \cdot)$.

(ii) With probability

$$\min\left\{1, \frac{\tilde{U}_k}{U_{k-1}} \cdot \frac{q(\tilde{X}_k, X_{k-1})}{q(X_{k-1}, \tilde{X}_k)}\right\}$$

accept and set $(X_k, U_k, \zeta_k) = (\tilde{X}_k, \tilde{U}_k, \tilde{\zeta}_k)$; otherwise reject and set $(X_k, U_k, \zeta_k) = (X_{k-1}, U_{k-1}, \zeta_{k-1})$.

If the Markov chain $(X_k, U_k, \zeta_k)_{k\geq 1}$ defined in Algorithm 1 is Harris-recurrent, then $n^{-1}\sum_{k=1}^n \zeta_k$ converges to $\pi(f^*)$ almost surely, because Algorithm 1 is easily checked to be a Metropolis-Hastings algorithm with target probability

$$\check{\pi}(\mathrm{d}x, \mathrm{d}u, \mathrm{d}\xi) := \frac{Q\big(x, (\mathrm{d}u, \mathrm{d}\xi)\big) u \mathrm{d}x}{\int \varpi_u(z)\mathrm{d}z}, \quad \text{satisfying} \quad \int \check{\pi}(\mathrm{d}x, \mathrm{d}u, \mathrm{d}\xi)\xi = \varpi(f).$$

The analogy of Assumption 6 and Assumption 2 is obvious.

The connections of pseudo-marginal MCMC and proper weighting suggests that IS type correction may sometimes be applied alternative to the pseudo-marginal method. This connection is particularly clear in the context of delayed acceptance (DA) MCMC [16, 60], which is a multiple-stage acceptance modification of the Metropolis-Hastings algorithm aiming for computational savings, and often based on approximate target distributions.

Our IS correction scheme has a direct connection with the following two-stage pseudo-marginal DA scheme based on an approximation π_a of the true probability density π. Such DA scheme with proposal densities $\{q(x, \cdot)\}_{x\in\mathsf{X}}$ as in Algorithm 1 can be summarised as follows:

Algorithm 2. Suppose Assumption 6 holds with $U = W$ and $\varpi(x) = w_u(x)$, $X_0 \in \operatorname{supp}(\pi)$ and for $k \geq 1$,

(i) Generate a proposal $\tilde{X}_k \sim q(X_{k-1}, \cdot)$
(ii) With probability

$$\min\left\{1, \frac{\pi_a(\tilde{X}_k)q(\tilde{X}_k, X_{k-1})}{\pi_a(X_{k-1})q(X_{k-1}, \tilde{X}_k)}.\right\}$$

continue to step (iii), otherwise set $(X_k, W_k, \zeta_k) = (X_{k-1}, W_{k-1}, \zeta_{k-1})$.
(iii) Generate $(\tilde{W}_k, \tilde{\zeta}_k) \sim Q(\tilde{X}_k), \cdot)$, and with probability $\min\{1, \tilde{W}_k/\tilde{W}_{k-1}\}$ set $(X_k, W_k, \zeta_k) = (\tilde{X}_k, \tilde{W}_k, \tilde{\zeta}_k)$; otherwise set $(X_k, W_k, \zeta_k) = (X_{k-1}, W_{k-1}, \zeta_{k-1})$.

The steps (i) and (ii) alone would implement a regular Metropolis-Hastings step targeting π_a, whereas the second-stage acceptance step (iii) alone would implement the accept-reject step of pseudo-marginal independence Metropolis-Hastings algorithm with proposal π_a. Our approach replaces the latter by weighting.

The rationale behind the DA scheme is based on the same two assumptions which make the IS corrected MCMC appealing. The evaluation of π_a is cheap compared with constructing the estimators $(\tilde{W}_k, \tilde{\zeta}_k) \sim Q(\tilde{X}_k)$, and π_a should approximate π reasonably well. The former directly affects the expected cost of the algorithm, by reducing the need of expensive estimators, and the latter assumption guarantees that the second-stage acceptance rate is typically close to one. Even though the DA version is guaranteed to be worse than direct Metropolis-Hastings targeting π in the Peskun sense [8], the lower cost often outweighs the difference by allowing a greater number of samples to be drawn with the DA scheme. The efficiency of DA schemes was studied by Banterle, Grazian, Lee and Robert [8]. Sherlock, Thiery and Golightly [81] studied the efficiency of pseudo-marginal with DA as in Algorithm 2, with diffusion limit techniques.

Contrary to the IS scheme, the noisy importance weights affect the DA Markov chain convergence, and may lead to slower mixing. For instance, if the distribution of $\tilde{W}_k$ in Algorithm 2 is unbounded, the chain exhibits 'sticky' behaviour, and is sub-geometrically ergodic [3]. This means that a central limit theorem fails for certain functions. In case of IS correction, the effect of noisy IS weights does not affect the MCMC chain, but their effect is more transparent in the expressions of asymptotic variances (Theorems 8 and 14).

7. General state space models and Monte Carlo inference

State space model (SSM) techniques are applied in a wide variety of applications [cf. 13]. We assume the existence of a family of SSMs indexed by parameter θ taking values on a space Θ. The models consist of a sequence of latent (unobserved) states $\alpha_{1:T}$ taking values in S_α, and observations $y_{1:T}$ taking values in Y. More specifically, the latent states $\alpha_{1:T}$ form a Markov chain with initial density $\mu_1^{(\theta)}(\alpha_1)$ and state transition densities $\mu_t^{(\theta)}(\alpha_t \mid \alpha_{t-1})$ for $t = 2, \ldots, T$. The observations $y_{1:T}$ are conditionally independent given $\alpha_{1:T}$ and follow the observation model $y_t \sim g_t^{(\theta)}(\cdot \mid \alpha_t)$. This model defines the following probability density on $\mathsf{S}_\alpha^T \times \mathsf{Y}^T$:

$$p^{(\theta)}(\alpha_{1:T}, y_{1:T}) = \prod_{t=1}^{T} g_t^{(\theta)}(y_t \mid \alpha_t)\mu_t^{(\theta)}(\alpha_t \mid \alpha_{t-1}),$$

where, by convention, $\mu_1^{(\theta)}(\alpha_1 \mid \alpha_0) := \mu_1^{(\theta)}(\alpha_1)$. The conditional distribution of $\alpha_{1:T}$ given $y_{1:T}$, often referred to as the smoothing distribution, has the form

$$p^{(\theta)}(\alpha_{1:T} \mid y_{1:T}) = p^{(\theta)}(\alpha_{1:T}, y_{1:T})/p^{(\theta)}(y_{1:T}),$$

where $p^{(\theta)}(y_{1:T}) := \int p^{(\theta)}(\tilde{\alpha}_{1:T}, y_{1:T})\mathrm{d}\tilde{\alpha}_{1:T}$ is a normalising constant.

We consider full Bayesian inference, where the parameters θ are equipped with a prior density $\mathrm{pr}(\theta)$, and the inference is based on the posterior distribution

$$p(\theta, \alpha_{1:T} \mid y_{1:T}) = \mathrm{pr}(\theta)p^{(\theta)}(\alpha_{1:T}, y_{1:T})/p(y_{1:T}).$$

with overall normalising constant $p(y_{1:T}) := \int \mathrm{pr}(\tilde{\theta})p^{(\tilde{\theta})}(y_{1:T})\mathrm{d}\tilde{\theta}$. We may apply our generic framework by taking $x = \theta$ and $z = \alpha_{1:T}$. This section is devoted to general techniques to generate, for given θ, random variables $(U, V^{(1:m)}, Z^{(1:m)})$ which satisfy

$$\mathbb{E}[U\xi(h)] = \int p^{(\theta)}(\alpha_{1:T}, y_{1:T})h(\alpha_{1:T})\mathrm{d}\alpha_{1:T} \quad \text{with} \quad \xi(h) := \sum_{i=1}^{m} V^{(i)}h(Z^{(i)}), \tag{10}$$

for all functions of interest $h : \mathsf{S}_\alpha^T \to \mathbb{R}$. Under Assumption 1, where $X_k = \theta_k$, we may construct a proper sampling scheme $(W_k, \xi_k)_{k\geq 1}$ by defining $U_k = \frac{\mathrm{pr}(\theta_k)}{\pi_a(\theta_k)}U_k$, where (U_k, ξ_k) satisfy (10) with $\theta = \theta_k$.

If $q^{(\theta)}(\alpha_{1:T} \mid y_{1:T})$ are probability densities such that

$$p^{(\theta)}(\alpha_{1:T}, y_{1:T}) > 0 \implies q^{(\theta)}(\alpha_{1:T} \mid y_{1:T}) > 0,$$

we may construct $(U, V^{(1:m)}, Z^{(1:m)})$ using importance sampling, by drawing independent $\alpha_{1:T}^{(i)} \sim q^{(\theta)}(\,\cdot \mid y_{1:T})$, setting $\omega^{(i)} := p^{(\theta)}(\alpha_{1:T}^{(i)}, y_{1:T})/q^{(\theta)}(\alpha_{1:T}^{(i)})$ for $i \in \{1{:}m\}$, and

$$U := \frac{1}{m}\sum_{i=1}^{m}\omega^{(i)}, \qquad Z^{(i)} := \alpha_{1:T}^{(i)}, \qquad \text{and} \qquad V^{(i)} := \begin{cases} m^{-1}\omega^{(i)}/U, & U > 0 \\ m^{-1}, & U = 0. \end{cases}$$

Note that (10) is satisfied for all integrable h, so $\mathcal{L} = L^1(\pi^*)$. It is often useful to combine such schemes as in Proposition 7. This allows, for instance, using pairs of antithetic variables as in [29], which can lead to substantial variance reduction.

Let us consider then a particle filter (PF) algorithm [37]; see also the monographs [13, 17, 25]. The algorithm is based on sequential approximation of the smoothing distribution using an interacting particle system. We shall first consider a generic version of the algorithm, using the following standard general ingredients [cf. 17].

(i) Proposal distributions: M_1 is a probability density on S_α and $M_t(\,\cdot \mid \alpha_{1:t-1})$ defines a density on S_α conditional on $\alpha_{1:t-1}$.
(ii) Potential functions: $G_t : \mathsf{S}_\alpha^t \to \mathbb{R}_+$.
(iii) Resampling laws: $\mathrm{Res}(\,\cdot \mid \bar{\omega}^{(1:m)})$ defines a probability distribution on $\{1{:}m\}^m$ for every discrete probability mass $\bar{\omega}^{(1:m)}$.

We then describe the PF algorithm, with the convention that whenever the index 'i' appears, it takes values in $1{:}m$.

Algorithm 3 (Particle filter)**.** Initial state:

(i) Sample $\alpha_1^{(i)} \overset{\text{i.i.d.}}{\sim} M_1^{(\theta)}(\,\cdot\,)$, set $\bar{\alpha}_1^{(i)} = \alpha_1^{(i)}$ and let $\omega_1^{(i)} := G_1(\alpha_1)$.

For $t = 2, \ldots, T$, do

(ii) Sample $A_{t-1}^{(1:m)} \sim \mathrm{Res}(\cdot \mid \bar{\omega}_{t-1}^{(1:m)})$, where

$$\bar{\omega}_{t-1}^{(i)} := \begin{cases} \omega_{t-1}^{(i)}/\omega_{t-1}^{*}, & \omega_{t-1}^{*} > 0 \\ m^{-1}, & \omega_{t-1}^{*} = 0, \end{cases} \quad \text{with} \quad \omega_{t-1}^{*} := \sum_{j=1}^{m} \omega_{t-1}^{(j)}.$$

(iii) Sample $\alpha_t^{(i)} \sim M_t(\cdot \mid \bar{\alpha}_{t-1}^{(A_{t-1}^{(i)})})$ and set $\bar{\alpha}_t^{(i)} = (\bar{\alpha}_{t-1}^{(A_{t-1}^{(i)})}, \alpha_t^{(i)})$.

(iv) Calculate $\omega_t^{(i)} := G_t(\bar{\alpha}_{t-1}^{(A_{t-1}^{(i)})}, \alpha_t)$.

Remark 16. If $\omega_t^* = 0$ in Algorithm 3, the algorithm may be terminated immediately — then all the estimators considered in Proposition 17 below equal zero. We prefer to avoid such explicit stopping mechanism here for mathematical convenience.

The following unbiasedness assumption allows the common multinomial, stratified, residual and systematic resampling methods [cf. 23]. Further generalisations of the resampling scheme, which allow for example the auxiliary PF [71], or so-called adaptive resampling schemes and intermediate algorithms between SIS and PF [e.g. 93], are not pursued here.

Assumption 7. The resampling procedure satisfies

$$\mathbb{E}\left[\sum_{i=1}^{m} \mathbf{1}\left(A^{(i)} = j\right)\right] = m\bar{\omega}^{(j)} \qquad \text{where} \qquad A^{(1:m)} \sim \mathrm{Res}(\cdot \mid \bar{\omega}^{(1:m)}),$$

for any $j \in \{1{:}m\}$ and any probability mass vector $\bar{\omega}^{(1:m)}$.

In order to apply the method in the SSM domain, we require the following consistency condition, which we discuss below Proposition 17.

Assumption 8. Suppose that

$$\int h(\alpha_{1:T}) \prod_{t=1}^{T} M_t(\alpha_t \mid \alpha_{1:t-1}) G_t(\alpha_{1:t}) \mathrm{d}\alpha_{1:T} = \int p^{(\theta)}(\alpha_{1:T}, y_{1:T}) h(\alpha_{1:T}) \mathrm{d}\alpha_{1:T},$$

for all $h : \mathsf{S}_\alpha^T \to \mathbb{R}$ with the latter integral well-defined and finite.

The following result, which is well-known at least for particle filter experts, summarises alternative ways how the random variables $(U, V^{(1:m)}, Z^{(1:m)})$ may be constructed from the PF output, in order to satisfy (10). Because the results stated below are scattered in the literature, and some are stated under different conditions, we provide a self-contained and concise proof of Proposition 17 in Appendix D.

Proposition 17. *Suppose that Assumptions 7 and 8 hold, and let $h : \mathsf{S}_\alpha^T \to \mathbb{R}$ be such that the integral in* (10) *is well-defined and finite. Consider the random variables generated in Algorithm 3, and let $U := \prod_{t=1}^{T} \left(\frac{1}{m}\omega_t^*\right)$.*

(i) The random variables $(U, V^{(1:m)}, Z^{(1:m)})$ where $V^{(i)} = \bar{\omega}_T^{(i)}$ and $Z^{(i)} = \bar{\alpha}_T^{(i)}$ satisfy (10).

Suppose then that $M_t(\alpha_t \mid \alpha_{1:t-1})G_t(\alpha_{1:t}) = C_t(\alpha_{t-1:t})$ *for all* $t \in \{1{:}T\}$ *and* $\alpha_{1:T} \in \mathsf{S}_\alpha^T$. *Define for* $t = 2, \dots, T$, *and any* $i_t, i_{t-1} \in \{1{:}m\}$, *the backwards sampling probabilities*

$$b_{t-1}(i_{t-1} \mid i_t) := \frac{\bar{\omega}_{t-1}^{(i_{t-1})} C_t(\alpha_{t-1}^{(i_{t-1})}, \alpha_t^{(i_t)})}{\sum_{\ell=1}^m \bar{\omega}_{t-1}^{(\ell)} C_t(\alpha_{t-1}^{(\ell)}, \alpha_t^{(i_t)})}, \quad \text{and} \quad b_T(i_T \mid i_{T+1}) = b_T(i_T) = \bar{\omega}_T^{(i_T)}.$$

(ii) Let $I_{1:T}$ *be a random variable generated recursively by* $I_T \sim b_T$, *and* $I_t \sim b_t(\cdot \mid I_{t+1})$ *for* $t = T-1, \dots, 1$. *The random variables* $(U, Z^{(1)}, V^{(1)})$ *satisfy* (10), *where* $V^{(1)} = 1$ *and* $Z^{(1)} = \alpha_{1:T}^{(I_{1:T})}$.

(iii) If $h(\alpha_{1:T}) = \hat{h}(\alpha_{t-1}, \alpha_t)$ *for some* $t \in \{2{:}T\}$, *that is,* h *is constant in all coordinates except* $t-1$ *and* t, *then, the random variables* $(U, Z^{(1:m,1:m)}, V^{(1:m,1:m)})$ *satisfy* (10) *(with* $\hat{h}$ *on the left), where*
(a) $Z^{(i,j)} := (\alpha_{t-1}^{(i)}, \alpha_t^{(j)})$,
(b) $V^{(i,j)} := b_{t-1}(i \mid j)\hat{\omega}_t^{(j)}$, *and where*
(c) $\hat{\omega}_T^{(i)} := \bar{\omega}_T^{(i)}$ *and* $\hat{\omega}_t^{(i)} := \sum_{k=1}^m \hat{\omega}_{t+1}^{(k)} b_t(i \mid k)$ *for* $t = T-1, \dots, t$.

(iv) If $h(\alpha_{1:T}) = \hat{h}(\alpha_t)$ *for some* $t \in \{1{:}T\}$, *then the random variables* $(V^{(1:m)}, Z^{(1:m)}, U)$ *satisfy* (10) *(with* $\hat{h}$ *on the left), where* $Z^{(i)} = \alpha_t^{(i)}$ *and* $V^{(i)} = \hat{\omega}_t^{(i)}$ *are defined in* (iiic).

Proposition 17 (i) corresponds to so-called filter-smoother [52]. This property was shown in [17, Theorem 7.4.2], although only in case of multinomial resampling. The same property is behind the original version of the particle marginal Metropolis-Hastings [1]. The statement holds also when the PF is applied with a general sequence of distributions $\pi_{1:t}^{(\theta)}$ rather than the SSM [17].

Proposition 17 (ii) corresponds to so-called backwards simulation smoothing [36], which has been developed recently within the conditional sequential Monte Carlo context [15, 59, 90]. This allows us to draw any number of independent trajectories (indices), and average them, in order to satisfy (10).

Proposition (iii) and its special case (iv) correspond to the forward-backward smoother [26]; see also [13]. It can be seen as Rao-Blackwellised version of (ii), but is applicable only when h has a specific form. This scheme can lead to lower variance, but it also has square complexity in m, rendering it inapplicable with large m.

The consistency Assumption 8 holds with commonly used potentials

$$G_t(\alpha_{1:t}) := \frac{g_t^{(\theta)}(y_t \mid \alpha_t)\mu_t^{(\theta)}(\alpha_t \mid \alpha_{t-1})}{M_t(\alpha_t \mid \alpha_{1:t-1})}, \tag{11}$$

assuming the following well-known minimal support condition

$$\pi_t^{(\theta)}(\alpha_{1:t}) > 0 \implies \pi_{t-1}^{(\theta)}(\alpha_{1:t-1})M_t(\alpha_t \mid \alpha_{1:t-1}) > 0, \qquad \text{for all } t \in \{1{:}T\},$$

$$\pi_t^{(\theta)}(\alpha_{1:t}) := \prod_{\ell=1}^t g_\ell^{(\theta)}(y_\ell \mid \alpha_\ell)\mu_\ell^{(\theta)}(\alpha_\ell \mid \alpha_{\ell-1}).$$

This choice of G_t satisfies $M_t(\alpha_t, \alpha_{1:t-1})G_t(\alpha_{1:t}) = g_t^{(\theta)}(y_t \mid \alpha_t)\mu_t(\alpha_t \mid \alpha_{t-1}) = C_t(\alpha_{t-1:t})$, and the expression of b_{t-1} simplifies to having $\mu_t(\alpha_t \mid \alpha_{t-1})$ in place of $C_t(\alpha_{t-1:t})$. The choice $M_t(\alpha_t \mid \alpha_{1:t-1}) = \mu_t^{(\theta)}(\alpha_t \mid \alpha_{t-1})$ is known as the bootstrap

filter following [37]. We discuss another choice of M_t and G_t, which is of not this form, but satisfies Assumption 8, in Section 8. It allows to take full advantage of approximations of the smoothing distributions $q^{(\theta)}(\alpha_{1:T} \mid y_{1:T})$.

We next formally state how Proposition 17 allows to use Algorithm 3 to derive a proper weighting scheme.

Corollary 18. *Let $(\theta_k)_{k\geq 1}$ be a Markov chain which is Harris ergodic with respect to π_a. Suppose each $(U_k, V_k^{(1:m)}, Z_k^{(1:m)})$ corresponds to an independent run of Algorithm 3 with $\theta = \theta_k$, as defined in Proposition 17* (i), (ii), (iii) *or* (iv). *Then, $(W_k, V_k^{(1:m)}, Z_k^{(1:m)})_{k\geq 1}$ with $W_k := \mathrm{pr}(\theta_k)U_k/\pi_a(\theta_k)$ provide a proper weighting scheme for target distribution $\pi^*(\theta, \alpha_{1:T}) = p(\theta, \alpha_{1:T} \mid y_{1:T})$ (Assumption 3), for the following classes of functions, respectively:*

(i) $\mathcal{L} = L^1(\pi^*)$, (iii) $\mathcal{L} = \{f \in L^1(\pi^*) : f(\theta, \alpha_{1:T}) = \hat{f}(\theta, \alpha_{t-1:t})\}$,

(ii) $\mathcal{L} = L^1(\pi^*)$, (iv) $\mathcal{L} = \{f \in L^1(\pi^*) : f(\theta, \alpha_{1:T}) = \hat{f}(\theta, \alpha_t)\}$.

Remark 19. The latter two cases in Corollary 18 are stated for a single marginal (pair), but it is clear that we may calculate estimates simultaneously for several marginal (pairs), so that Proposition 17 (iii) is applicable for every function which is of the form $f\sum_{t=1}^{T-1} f_t(\theta, \alpha_{t:t+1})$ and Proposition 17 (iv) whenever $\sum_{t=1}^{T} f_t(\theta, \alpha_t)$. See also the general discussion of smoothing functionals in [13, §4.1.2].

We next discuss briefly one implication of Proposition 17 outside the main focus of this paper, which is widely known for particle filtering experts. Namely, in the general state space smoothing context, that is, where $\theta \in \Theta$ is constant, Proposition 17 suggests easily parallelisable particle smoothing algorithms, which also admit easily calculated consistent confidence interval estimates.

We state formally the following result, whose proof follows similarly as the proof of Theorem 9 in Appendix B.

Proposition 20. *Suppose $\theta \in \Theta$ is fixed, and let $(Z_k^{(1:m)}, V_k^{(1:m)}, U_k)_{k\geq 1}$ correspond to independent realisations of random variables defined in Proposition 17.*

(i) If the corresponding conditions of Proposition 17 are satisfied, then the estimator

$$E_n(h) := \sum_{k=1}^{n} \tilde{U}_{k,n}\xi_k(h) \xrightarrow[\text{a.s.}]{n\to\infty} \mu_h := \int p^{(\theta)}(\alpha_{1:T} \mid y_{1:T})h(\alpha_{1:T})\mathrm{d}\alpha_{1:T},$$

where $\tilde{U}_{k,n} := U_k/\sum_{j=1}^{n} U_j$.

(ii) If also $\sigma_*^2 := \mathbb{E}\big[U_1^2(\xi_k(h) - \mu_h)^2\big] < \infty$, then

$$\sqrt{n}\big[E_n(h) - \mu_h\big] \xrightarrow[\mathrm{d}]{n\to\infty} N(0, \sigma^2), \qquad \textit{where} \qquad \sigma^2 := \frac{\sigma_*^2}{p^{(\theta)}(y_{1:T})^2}.$$

(iii) If in addition $\mathbb{E}\big[U_1^2\big] < \infty$, then $nv_n \to \sigma^2$, almost surely, where

$$v_n := \sum_{k=1}^{n} \tilde{U}_{k,n}^2\big(\xi_k(h) - E_n(h)\big)^2.$$

Proposition 20 formalises the conditions under which the confidence intervals $\left[E_n(f) \pm \beta\sqrt{v_n}\right]$ are asymptotically consistent, where β corresponds to the desirable standard Gaussian quantile. Calculation of consistent confidence intervals for a single realisation of a particle smoothing algorithm is not a straightforward task, but there is some recent progress [55]. Another recently suggested method [44] relies on unbiased estimators obtained by coupling of conditional sequential Monte Carlo and randomisation as in [66, 75].

8. State space models with linear-Gaussian state dynamics

We now consider a special case of the general state space model introduced in Section 7, where both S_α and Y are Euclidean, the state dynamics $\mu_t^{(\theta)}(\alpha_t \mid \alpha_{t-1})$ are linear-Gaussian, but the observational distribution $g_t^{(\theta)}(y_t \mid \alpha_t)$ may be non-Gaussian and/or non-linear. More specifically, $\mu_1 = N(a_1, P_1)$ and

$$\alpha_{t+1} := F_t\alpha_t + \eta_t, \qquad \eta_t \overset{\text{i.i.d.}}{\sim} N(0, Q_t), \qquad t = 1, 2, \ldots, T, \tag{12}$$

where a_1, P_1, F_t, Q_t are parameters. Our approach allows, for instance, exponential family observation models with Gaussian, Poisson, binomial, negative binomial, and Gamma distributions, and a stochastic volatility model.

Compared to the general state space form, this class of models may seem restrictive, but it still contains a large number of commonly used statistical models, such as structural time series models, cubic splines, generalised linear mixed models, and autoregressive integrated moving average (ARIMA) models. The benefit of working within this subset of general state space models is that we have access to approximations taking advantage of the specific structure of the model.

8.1. Marginal approximation.

Our approximation scheme is based on [28, 80], and relies on a Laplace approximation $\tilde{p}^{(\theta)}(\alpha_{1:T} \mid y_{1:T})$ of the conditional densities $p^{(\theta)}(\alpha_{1:T} \mid y_{1:T})$. The approximating Gaussian model is found by using an iterative process involving repeated runs of Kalman filter and smoother, and is closely related to iterative reweighted least squares algorithm used in generalised linear models context [65]. In the approximating Gaussian model, the observation equation is replaced by a linear-Gaussian one

$$\tilde{y}_t = H_t\alpha_t + \epsilon_t, \qquad \epsilon_t \overset{\text{i.i.d.}}{\sim} N(0, R_t), \tag{13}$$

where H_t is defined in $g_t^{(\theta)}(y_t|\alpha_t)$, and the pseudo-observations $\tilde{y}_t$ and the variances R_t are constructed using the first and second derivatives of $\log g^{(\theta)}(y_t \mid \alpha_t)$ with respect to α_t at $\hat{\alpha}_t$, the conditional mode estimate of the approximate model [29].

Consider the following decomposition of the marginal likelihood:

$$p^{(\theta)}(y_{1:T}) = \tilde{p}^{(\theta)}(\tilde{y}_{1:T}) \frac{g^{(\theta)}(y_{1:T} \mid \hat{\alpha}_{1:T})}{\tilde{g}^{(\theta)}(\tilde{y}_{1:T} \mid \hat{\alpha}_{1:T})} \mathbb{E}\left[\frac{g^{(\theta)}(y_{1:T} \mid \alpha_{1:T})/g^{(\theta)}(y_{1:T} \mid \hat{\alpha}_{1:T})}{\tilde{g}^{(\theta)}(\tilde{y}_{1:T} \mid \alpha_{1:T})/\tilde{g}^{(\theta)}(\tilde{y}_{1:T} \mid \hat{\alpha}_{1:T})}\right], \tag{14}$$

where $g^{(\theta)}(y_{1:T} \mid \alpha_{1:T}) := \prod_{t=1}^T g_t^{(\theta)}(y_t \mid \alpha_t)$, $\tilde{g}^{(\theta)}(\tilde{y}_{1:T} \mid \alpha_{1:T}) := \prod_{t=1}^T \tilde{g}_t^{(\theta)}(\tilde{y}_t \mid \alpha_t)$, where $\tilde{g}_t^{(\theta)}(\tilde{y}_t \mid \alpha_t)$ correspond to (13), the expectation is taken with respect to $\tilde{p}^{(\theta)}(\alpha_{1:T} \mid \tilde{y}_{1:T})$, and $\tilde{p}^{(\theta)}(\tilde{y}_{1:T})$ corresponds to the marginal likelihood of the pseudo-observations, which may be calculated using the Kalman filter.

If the pseudo-likelihoods $\tilde{g}_t^{(\theta)}$ are nearly proportional to the true likelihoods $g_t^{(\theta)}$ in high-probability regions of $\tilde{p}^{(\theta)}(\alpha_{1:T} \mid y_{1:T})$, the expectation in (14) is close to

one. We base our approximation of the marginal density π_a to (14), dropping the expectation, that is,

$$\pi_a(\theta) \propto \mathrm{pr}(\theta)\tilde{p}^{(\theta)}(y_{1:T})\frac{g^{(\theta)}(y_{1:T} \mid \hat{\alpha}_{1:T})}{\tilde{g}^{(\theta)}(y_{1:T} \mid \hat{\alpha}_{1:T})}. \tag{15}$$

This approximation defines, in fact, a joint approximate distribution $\pi_a(\theta, \alpha_{1:T}) = \pi_a(\theta)\tilde{\pi}^{(\theta)}(\alpha_{1:T} \mid y_{1:T})$.

The approximate marginal $\pi_a(\theta$ above is based on a Laplace approximation of the smoothing distribution, which needs to be reconstructed for each new value of θ. Let us call this approach local approximation, and consider a faster, but less accurate, variant where pseudo-observations and variances are constructed only once, at the maximum likelihood estimate of θ. We call such an approach the global approximation **G**.

8.2. **Proper weighting schemes.** Section 7 describes generic methods which lead to proper weighting schemes. The first approach is obvious, using the approximate smoothing distributions $\tilde{\pi}^{(\theta)}(\alpha_{1:T} \mid y_{1:T})$ as IS proposals. The simulation smoother algorithm of [30] allows to use antithetic variables in order to reduce the variance. We call this scheme **SPDK** following [80].

In case of the particle filter, there is no obvious choice of parameters. The off-the-shelf solution is the bootstrap filter [37], which we abbreviate **BSF**. Because the BSF does not use the approximation $\tilde{\pi}^{(\theta)}(\alpha_{1:T} \mid y_{1:T})$ at all, it may be wasteful, and may lead to inefficient estimators. If we consider the standard setting (11), it is commonly recognised that the 'optimal' choice is to set $M_t(\alpha_t \mid \alpha_{1:t-1}) \propto g_t^{(\theta)}(y_t \mid \alpha_t)\mu_t^{(\theta)}(\alpha_t \mid \alpha_{t-1})$ [cf. 13]. However, the developments in [39, 92] suggest that other choices which are not of the form (11) may be more efficient. In the present context, the following choice

$$M_t(\alpha_t \mid \alpha_{1:t-1}) = \tilde{p}^{(\theta)}(\alpha_t \mid \alpha_{t-1}, y_{1:T}), \quad \text{and} \quad G_t(\alpha_{1:t}) = \frac{g_t^{(\theta)}(y_t \mid \alpha_t)}{\tilde{g}_t^{(\theta)}(\tilde{y}_t \mid \alpha_t)}$$

is natural, and would be optimal in our setting if G_t were constants. As they are often approximately so, we use this choice of M_t and G_t and call this method ψ-**PF**, following the terminology of [39].

9. Experiments

We now turn into an illustration of our generic framework, where we applied the approximations in Section 8 with two specific models. We compared several inference algorithms in our experiments, using the following five base methods:

- **aMCMC:** Approximate MCMC targeting $\pi_a(\theta)$, and for each accepted θ, sampling one realisation from $\tilde{p}^{(\theta)}(\alpha_{1:T} \mid y_{1:T})$.
- **PM:** Pseudo-marginal MCMC (Algorithm 1).
- **DA:** Delayed acceptance pseudo-marginal MCMC (Algorithm 2).
- **IS1:** Simple importance sampling estimator given in (4) with m samples.
- **IS2:** Jump chain estimator as in (7), with IS estimators based on m samples, independent of N_i.

The last four algorithms are exact, in the sense that they admit correct limiting properties. We experimented the exact methods with three different IS correction schemes, SPDK, BSF and ψ-PF, described in Section 8.2. For BSF and ψ-PF, the filter-smoother estimates as in Proposition 17 (i) were used.

We used a random walk Metropolis algorithm with a Gaussian proposal distribution, whose covariance was adapted during burn-in using the algorithm described in [88]. The target mean acceptance rate was set to 0.234, but instead of the total acceptance rate, the acceptance rate based on the Gaussian approximation was optimised, except for PM-BSF where the total acceptance rate was optimised, as Gaussian approximation was unavailable. The maximum likelihood estimate of θ was always used as the starting value of MCMC.

The following algorithmic modifications were also experimented:

Global approximation: If the global Laplace approximation as discussed in Section 8.1 is used, it is indicated in parenthesis such as aMCMC(G).

Parallel computation: In tests where we apply multiple CPUs, the number of cores is shown in parenthesis such as IS2(8).

Sub-sampling: In most of our experiments, we used sub-sampling as in Proposition 6, and sampled one trajectory $\alpha_{1:T}$ per each accepted state.

All the experiments were conducted in `R` [74] using the `bssm` package which is available online [42]. The experiments were run on a Linux server with eight octa-core Intel Xeon E7-8837 2.67GHz processors with total 1TB of RAM.

The average computational complexity per iteration, or cost, of the exact methods, when parallel implementation aspects are not considered, are typically ordered such that the IS2 costs less than DA, and IS1 and PM are comparable, having a higher cost than DA. In particular, DA typically requires fewer simulations than IS1 and PM, which both require simulation at every iteration. The cost of DA is related to the mean second stage acceptance rate. In the extreme case of rate one, the DA typically requires the same number of computations on average as IS2. The most natural pairs of comparison are, therefore, PM against IS1 and DA against IS2.

For all the exact methods, we chose the IS accuracy parameter m based on a pilot experiment and the theoretical considerations in [27]. This choice is nearly optimal for pseudo-marginal MCMC, but we keep the same m for all methods, for comparability. More specifically, [27] suggest setting m so that the standard deviation of the logarithm of the likelihood estimate, denoted with δ, is around 1.2 in the neighbourhood of the posterior mean of θ. We note that [27] require the computational cost to be proportional to m, but here the cost is only asymptotically so, including a non-negligible initial overhead due to the approximation. Therefore, it is likely that for small m, say less than 100, the optimal choice is likely to be somewhat higher than what is suggested by the guidelines of [27]. Nevertheless, we disregard this issue in the following examples for simplicity.

In each experiment, we calculated the Monte Carlo estimators several times independently. The standard errors (SE) of the estimates, as well as the inverse relative efficiency (IRE) are reported. The IRE, defined as the mean square error (MSE) of the estimate multiplied by the average computation time, provides a justified way to compare Monte Carlo algorithms with different costs [35]. When

calculating the MSE, we used the average over all estimates from all unbiased algorithms as the ground truth.

9.1. **Poisson state space model.** The first model is of the following form:

$$\begin{aligned} g_t^{(\theta)}(y_t \mid \alpha_t) &= \text{Poisson}(y_t; \exp(\mu_t)), \\ \mu_{t+1} &= \mu_t + \nu_t + \sigma_\eta \eta_t, & \eta_t &\sim N(0,1), \\ \nu_{t+1} &= \nu_t + \sigma_\xi \xi_t, & \xi_t &\sim N(0,1), \end{aligned}$$

with $\alpha_1 \sim N(0, 0.1I)$. For testing our algorithms, we simulated a single set of observations $y_{1:100}$ from this model with α_1 fixed to $(0,0)$, and $\theta = (\sigma_\eta, \sigma_\xi)$ fixed to $(0.1, 0.01)$. We used a uniform prior $U(0, 2s)$ for the parameters, where the cut-off parameter s was set to 1.6 based on the sample standard deviation of $\log(y_{1:T})$, where zeros were replaced with 0.1. Results were not sensitive to this upper bound.

We ran pilot experiments with θ set to its posterior mean $(0.09, 0.02)$ in order to optimise m. With $m = 10$, the standard deviation δ was already around 0.1 for SPDK and ψ-PF. For BSF, $m = 200$ was sufficient to obtain the desired $\delta \approx 1.2$. For all algorithms, we used 100,000 MCMC iterations with the first half discarded as burn-in. We performed two experiments. In the first, we stored only one realisation of $\alpha_{1:T}$ at each iteration, and the resulting samples were used to calculate the posterior mean estimates of θ and $\alpha_{1:T}$. In the second experiment, we used all the realisations of states used in computing the likelihood estimate, and computed the posterior mean and variance estimates of $\alpha_{1:T}$. We ran all the algorithms independently 1000 times. The IREs were scaled so that all IREs of PM-BSF were fixed to one, for easier comparison.

Tables 1 and 2 summarise the results of the first experiment using aMCMC and BSF variants of the exact methods. The aMCMC took only a fraction of the time required by PM, making it tempting in terms of IRE in many cases. However, the bias of aMCMC may lead into problems, as demonstrated by a noticeable bias in the state variables. The overall acceptance rate of DA was around 0.104, and for all other methods the total acceptance rate was very close to the nominal 0.234.

The DA had slightly higher SEs than PM, but less than half of the computational cost. For IS1 and IS2, the SEs were slightly smaller than with PM and DA, respectively. Because the cost of IS2 was roughly half of DA, it outperformed DA by about a factor of two in terms of IRE. Similarly, IS1 was slightly faster than PM, and outperformed PM in terms of IREs by a clear margin.

There were no significant differences between the posterior mean estimates and the SEs from local and global approximation variants, but the global variant was about 45 seconds faster, leading to smaller IREs. Note that for PM-BSF no approximation was computed, so the results in Tables 1 and 2 for local and global approximation are identical for PM.

Table 3 summarises the results with SPDK. Compared to the previous results from Tables 1 and 2, the execution times of all the methods were much smaller due to the smaller number of samples needed for likelihood computation. Also because of small m, the speed-up from DA and IS1/IS2 over PM was limited, which is natural as the cost cannot be lower than with aMCMC. The total acceptance rate of DA-SPDK was also close to 0.234. The SEs were very close to each other,

TABLE 1. Results for Poisson model using aMCMC and BSF variants, local approximation and sub-sampling.

		aMCMC	PM	DA	IS1	IS2
time (s)		54	676	281	600	166
mean	σ_η	0.093	0.093	0.093	0.093	0.093
	σ_ξ	0.016	0.016	0.016	0.016	0.016
	μ_1	-0.064	-0.075	-0.075	-0.075	-0.075
	μ_{100}	2.629	2.618	2.618	2.618	2.619
SE	σ_η	2.0e-05	2.9e-05	3.8e-05	2.3e-05	3.4e-05
	σ_ξ	6.1e-06	8.4e-06	1.1e-05	6.5e-06	8.9e-06
	μ_1	1.1e-04	1.4e-04	2.1e-04	1.1e-04	1.8e-04
	μ_{100}	7.1e-05	8.9e-05	1.4e-04	7.4e-05	1.2e-04
IRE	σ_η	0.039	1.000	0.721	0.535	0.336
	σ_ξ	0.042	1.000	0.676	0.537	0.278
	μ_1	0.561	1.000	0.911	0.609	0.406
	μ_{100}	1.211	1.000	1.049	0.623	0.441

TABLE 2. Results for Poisson model using aMCMC and BSF variants, global approximation and sub-sampling.

		aMCMC	PM	DA	IS1	IS2
time (s)		11	676	235	549	120
mean	σ_η	0.093	0.093	0.093	0.093	0.093
	σ_ξ	0.016	0.016	0.016	0.016	0.016
	μ_1	-0.065	-0.075	-0.075	-0.075	-0.075
	μ_{100}	2.631	2.618	2.618	2.618	2.618
SE	σ_η	2.0e-05	2.9e-05	3.8e-05	2.2e-05	3.2e-05
	σ_ξ	6.2e-06	8.4e-06	1.1e-05	6.8e-06	8.8e-06
	μ_1	1.0e-04	1.4e-04	2.1e-04	1.1e-04	1.7e-04
	μ_{100}	6.7e-05	8.9e-05	1.3e-04	7.4e-05	1.2e-04
IRE	σ_η	0.012	1.000	0.596	0.476	0.218
	σ_ξ	0.052	1.000	0.564	0.530	0.197
	μ_1	0.085	1.000	0.779	0.527	0.273
	μ_{100}	0.333	1.000	0.804	0.563	0.305

and the differences between global and local approximation variants were again mostly due to cost. The lower cost and somewhat smaller SEs than with BSF led to about 20 times smaller IREs than in PM-BSF. The results for ψ-PF given in Table 4 are similar to SPDK, with slightly higher costs and IREs.

The results of the second experiment shown in Table 5 were similar to the first experiment, with somewhat higher costs and smaller SEs. The results with SPDK are omitted. In this case, the higher computational costs of IS1 were compensated by more accurate estimates, resulting typically smaller IREs than with IS2. The

TABLE 3. Results for Poisson model using SPDK variants and sub-sampling.

		Local approximation				Global approximation			
		PM	DA	IS1	IS2	PM	DA	IS1	IS2
time (s)		78	61	71	53	35	17	28	10
SE	σ_η	2.1e-05	2.1e-05	2.1e-05	2.1e-05	2.0e-05	2.1e-05	2.1e-05	2.1e-05
	σ_ξ	6.3e-06	6.4e-06	6.3e-06	6.3e-06	6.3e-06	6.5e-06	6.5e-06	6.5e-06
	μ_1	1.0e-04	1.1e-04	1.0e-04	1.0e-04	1.0e-04	1.1e-04	1.0e-04	1.1e-04
	μ_{100}	7.0e-05	7.2e-05	7.0e-05	7.2e-05	7.3e-05	7.0e-05	7.1e-05	7.2e-05
IRE	σ_η	0.060	0.047	0.056	0.042	0.025	0.013	0.022	0.008
	σ_ξ	0.064	0.052	0.059	0.044	0.029	0.015	0.025	0.009
	μ_1	0.063	0.055	0.057	0.042	0.027	0.016	0.023	0.009
	μ_{100}	0.072	0.059	0.067	0.052	0.034	0.016	0.027	0.010

TABLE 4. Results for Poisson model using ψ-PF variants and sub-sampling.

		Local approximation				Global approximation			
		PM	DA	IS1	IS2	PM	DA	IS1	IS2
time (s)		115	78	83	62	72	34	38	19
SE	σ_η	2.0e-05	2.2e-05	2.1e-05	2.1e-05	2.0e-05	2.1e-05	2.1e-05	2.1e-05
	σ_ξ	6.1e-06	6.4e-06	6.3e-06	6.3e-06	6.4e-06	6.5e-06	6.5e-06	6.5e-06
	μ_1	1.0e-04	1.1e-04	1.1e-04	1.1e-04	1.0e-04	1.1e-04	1.1e-04	1.0e-04
	μ_{100}	7.0e-05	7.2e-05	6.9e-05	7.2e-05	7.1e-05	7.3e-05	7.1e-05	7.3e-05
IRE	σ_η	0.082	0.068	0.065	0.049	0.052	0.028	0.030	0.015
	σ_ξ	0.091	0.068	0.069	0.051	0.061	0.031	0.034	0.017
	μ_1	0.097	0.071	0.076	0.053	0.056	0.030	0.033	0.015
	μ_{100}	0.106	0.075	0.074	0.060	0.068	0.034	0.036	0.019

TABLE 5. Results for Poisson model using BSF and ψ-PF variants, local approximation and without sub-sampling.

		BSF				ψ-PF			
		PM	DA	IS1	IS2	PM	DA	IS1	IS2
time (s)		900	509	819	215	132	91	91	63
SE	μ_1	1.1e-04	1.7e-04	4.6e-05	1.4e-04	5.4e-05	5.5e-05	2.9e-05	5.4e-05
	μ_{100}	2.4e-05	3.6e-05	1.8e-05	3.0e-05	2.9e-05	3.0e-05	1.9e-05	2.9e-05
IRE	μ_1	0.857	1.113	0.136	0.343	0.029	0.021	0.006	0.014
	μ_{100}	0.099	0.124	0.051	0.037	0.021	0.016	0.006	0.010

benefits from using all samples are limited particularly in case of μ_1, which is likely because of the nature of the filter-smoother estimates.

9.2. **Stochastic volatility model.** As a second illustration, we considered the daily returns for the S&P index from 4/1/1995 to 28/9/2016, with total number

of observations $T = 5473$. The data was analysed using the following stochastic volatility (SV) model

$$y_t = \exp(\alpha_t/2)\epsilon_t, \qquad \epsilon_t \overset{\text{i.i.d.}}{\sim} N(0,1)$$
$$\alpha_{t+1} = \mu + \phi(\alpha_t - \mu) + \sigma_\eta \eta_t, \qquad \eta_t \overset{\text{i.i.d.}}{\sim} N(0,1),$$

with $\alpha_1 \sim N(\mu, \sigma_\eta^2/(1-\phi^2))$. We used a uniform prior for ϕ on $[-0.9999, 0.9999]$. For σ_ϕ we used a half-Gaussian distribution with prior standard deviation 5, and we used Gaussian distribution for μ, also with standard deviation 5.

This experiment was much more challenging than the Poisson model. Sequential importance sampling was expected to be problematic, as it is well-known to deteriorate exponentially in T, unlike the particle filter which was expected to scale much better in T [91]. In addition, it is known that for this particular model, the importance weights may have a large variability [53, 72]. While in principle both SPDK and ψ-PF should be affected by such fluctuations, we did not observe any problems with ψ-PF in our experiments. We note that the large variability may be more problematic with PM and DA, leading into 'sticky' behaviour [3, 5].

As in the Poisson experiment, we chose m according to [27]. For ψ-PF, δ was approximately 1.1 using only $m = 10$ samples. For the other methods, we targeted the same $\delta \approx 1.1$, achieved with $m = 70$ in case of SPDK and $m = 3400$ in case of BSF. For all the algorithms, we used 100,000 MCMC iterations with the first half discarded as burn-in.

Table 6 gives results based on 100 independent realisations of the estimators using aMCMC and the BSF-variants based on local approximation. The PM and IS1 were omitted because of their high computational costs. Note that the IREs are scaled here with respect to DA, unlike in the Poisson experiment. The detailed results with global approximation are omitted, except for aMCMC which shows computational savings of slightly over one hour.

The aMCMC appeared to exhibit a noticeable bias, this time also for hyperparameters θ. The bias was higher with global approximation for μ, which may be because the SV model is more sensitive to the parameters than the Poisson model. Despite of the bias, the low cost of aMCMC makes it competitive in terms of IRE for many test functions, as in the Poisson experiment. There were no significant differences in the SEs between the exact methods when using the local and global approximation schemes.

The DA-BSF took about 46 hours to run, whereas IS2-BSF algorithm gave similar results in 25 hours. As the number of particles was high, we expected significant benefits from parallelisation. With 8 cores, the execution time dropped to 4.6 hours, which is very close to ideal. The total acceptance rate of DA was 0.1, and the differences of SEs for DA and IS2 were small.

Table 7 summarises the results with the SPDK and ψ-PF schemes. Due to a much smaller m, the DA-SPDK and DA-ψ-PF were an order of magnitude faster than DA-BSF, taking less than two hours to complete. Because of smaller m, a significant portion of the total execution time was due to the approximation. The acceptance rates for PM-SPDK and DA-SPDK were 0.09 and 0.06, respectively. Diagnostics from the individual runs sometimes showed poor mixing for these

TABLE 6. Results for the SV model using BSF and aMCMC variants.

		aMCMC	aMCMC(G)	DA	IS2	IS2(8)
time (h)		1.3	0.2	46.3	25.2	4.6
mean	ϕ	0.984	0.984	0.984	0.984	0.984
	σ_η	0.167	0.171	0.169	0.169	0.169
	μ	-0.190	-0.250	-0.191	-0.190	-0.190
	α_1	-1.780	-1.790	-1.749	-1.746	-1.744
	α_{5473}	-0.557	-0.572	-0.471	-0.473	-0.473
SE	ϕ	4.8e-06	4.9e-06	1.2e-05	9.0e-06	8.5e-06
	σ_η	2.0e-05	2.0e-05	3.5e-05	3.2e-05	3.4e-05
	μ	2.3e-04	2.5e-04	4.6e-04	3.3e-04	2.9e-04
	α_1	6.6e-04	6.7e-04	1.6e-03	1.2e-03	1.3e-03
	α_{5473}	6.2e-04	5.6e-04	1.0e-03	1.0e-03	1.1e-03
IRE	ϕ	0.083	0.062	1.000	0.304	0.050
	σ_η	0.726	0.298	1.000	0.483	0.096
	μ	0.008	0.747	1.000	0.287	0.042
	α_1	0.133	0.035	1.000	0.321	0.071
	α_{5473}	1.887	0.417	1.000	0.540	0.112

TABLE 7. Results for SV model using SPDK and ψ-PF variants.

		SPDK				ψ-PF			
		PM	DA	IS1	IS2	PM	DA	IS1	IS2
time (h)		4.4	1.9	2.8	1.5	2.4	1.4	1.5	1.3
SE	ϕ	3.9e-05	4.9e-05	4.1e-05	4.5e-05	7.3e-06	1.0e-05	6.2e-06	9.5e-06
	σ_η	2.0e-04	3.1e-04	1.1e-04	1.5e-04	3.1e-05	4.4e-05	2.4e-05	3.4e-05
	μ	1.6e-03	3.6e-03	9.1e-04	1.6e-03	3.3e-04	4.3e-04	2.6e-04	3.8e-04
	α_1	8.9e-03	1.1e-02	3.8e-03	5.3e-03	1.2e-03	1.4e-03	7.0e-04	1.3e-03
	α_{5473}	8.6e-03	7.0e-03	2.9e-03	5.2e-03	1.0e-03	1.5e-03	6.2e-04	1.2e-03
IRE	ϕ	1.015	0.696	0.684	0.483	0.021	0.024	0.009	0.017
	σ_η	3.090	3.307	0.603	0.710	0.044	0.055	0.016	0.028
	μ	1.208	2.544	0.228	0.404	0.026	0.027	0.010	0.020
	α_1	3.054	1.883	0.346	0.373	0.029	0.026	0.007	0.018
	α_{5473}	6.574	1.871	0.444	0.810	0.057	0.064	0.012	0.039

variants, and despite the large reductions in execution time, the IREs of PM-SPDK and DA-SPDK were worse than PM-BSF. We observed also cases with few very large correction weights in IS1 and IS2, which had an impact also on their efficiencies.

For PM and DA algorithms using ψ-PF, the total acceptance rates were 0.15 and 0.1 respectively, and the SEs of DA-ψ-PF were comparable with the DA-BSF. We did not experience problems with mixing or overly large weights with ψ-PF, which suggests that ψ-PF may be more robust than SPDK.

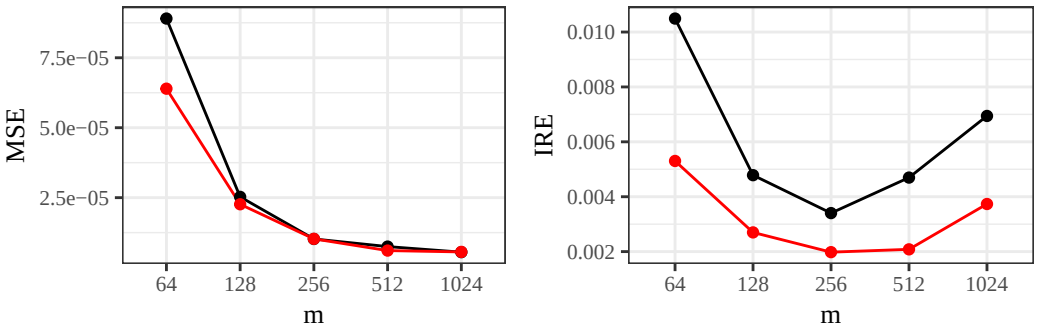

FIGURE 1. Average MSE and IRE of $(\sigma_\eta, \sigma_\xi, \mu_1, \mu_{100})$ for different values of m in the Poisson model with BSF. DA is shown in black and IS2 in red.

9.3. **Discussion on results.** Based on our results, IS1 and IS2 appear as competitive alternatives to PM and DA, respectively, with typically similar or slightly smaller SEs and costs. These differences were noticeable in case of BSF with moderate m, but smaller with SPDK and ψ-PF, where m is small. The global approximation gave additional performance boost in our experiments, without compromising the accuracy of the estimates.

Our experiments suggest that ψ-PF is generally preferable over SPDK, as it scales better with the length of the time series and appears less prone to fluctuations in weights. The BSF is easier to implement, and does not require approximation, but suffers from inferior performance. The IS2 appears to be competitive against DA even in terms of raw computational cost, and further significant performance benefits may be obtained with parallel hardware.

As noted earlier, the use of the guidelines by [27] are not necessary optimal in case of a significant cost due to the approximation. We did a small experiment about how the choice of m affects the IRE when using BSF in the Poisson model, and ψ-PF in the SV model. Figures 1 and 2 show the average MSE and IRE as a function of m using the DA and IS2 algorithms. In the Poisson-BSF case in Figure 1, the suggested value of $m = 200$ was nearly optimal, and both DA and IS2 behaved similarly in terms of m. The MSEs of IS2 were similar to DA, and the IREs were clearly smaller. In the SV-ψ-PF case shown in Figure 2, the optimal m for DA and IS2 was around 50, which was higher than the value chosen based on [27]. This is likely because of the non-negligible overhead cost of the approximation. The IS2 dominated DA again regardless of the value of m.

10. Discussion

We suggested to use MCMC targeting a marginal of an approximate posterior distribution based on a marginal likelihood approximation, and consequent IS type correction in order to perform consistent, efficient and easily parallelisable inference over the full posterior. We focused on state space models (SSMs), where sequential importance sampling and particle filters provide natural and efficient correction mechanisms. Our specific application was based on a Laplace approximation of SSMs with linear-Gaussian state dynamics, but arbitrary observation distributions.

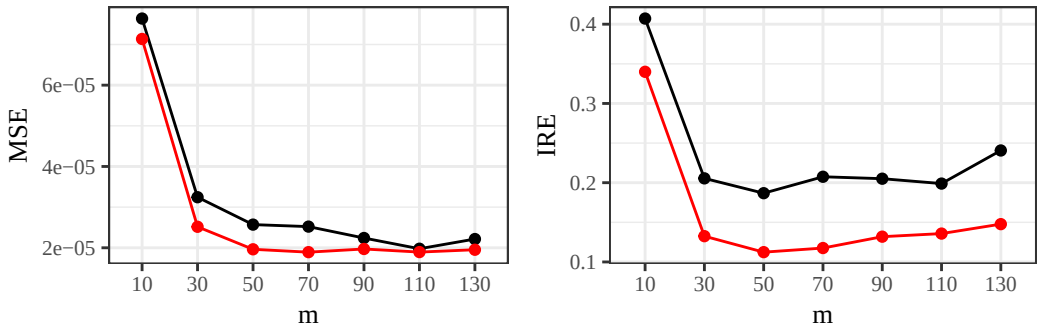

FIGURE 2. Average MSE and IRE of $(\phi, \sigma_\eta, \mu, \alpha_1, \alpha_{5473})$ for different values of m in the SV model with ψ-PF. DA is shown in black and IS2 in red.

We are currently investigating extensions of the Laplace approximations to a more general class of SSMs. Alternative approximation methods in the SSM context include simple linearisation-based methods such as the extended Kalman filter, Moment matching, or variational approximations; see the recent monograph [79]. Gaussian latent variable models beyond SSMs could be handled also with Laplace approximations, similar to [78]. We note that, unlike with purely approximate inference, the approximate posterior π_a need not be necessarily very accurate, like the performance of the global approximation in our experiment demonstrates.

We note that various particle filter variations, such as Rao-Blackwellised PF, and alternative resampling strategies, apply directly [13]. Randomised sequential quasi-Monte Carlo [32] can be useful, when applicable. Particle filters can also be useful beyond the state space context [18]. The correction could be based on exact sampling for diffusions [11], and debiasing tricks and multilevel Monte Carlo [33, 41, 66, 75, 89] may be applicable for instance when the true target distribution is based on a continuum model, and only discretisations are accessible.

Our experimental findings are promising, and we believe that our approach has a lot of potential for practical developments. Efficient use of approximations may enable exact inference with a large class of models where only approximate inference is currently feasible. Our approach may also be used as a verification tool, to ensure that the bias from an approximation, such as the integrated nested Laplace approximation [78], is controlled.

Importance sampling is known to be generally hard to implement efficiently in higher dimensions, and our setting is no exception. However, it may help that the approximation π_a is required only on a marginal distribution, which is often lower dimensional, like in our SSM applications. Indeed, in the SSM context our approach relies on similar assumptions which make particle MCMC [1] appealing, with the added assumption of the existence of suitable smoothing distribution approximations. If such approximations are available, our experiments suggest that our approach can often be preferable to particle MCMC or its delayed acceptance version.

One potential issue of importance sampling is when the importance weight $w_u(x) \propto \pi(x)/\pi_a(x)$ is unbounded. This can happen even if π_a closely resembles

π in the 'centre' of X, if X is unbounded. We suggest two possible modifications which may be used in order to ensure bounded w_u. The first is to use tempering, that is, if π_a is our original approximate marginal, the chain $(X_k)_{k\geq 1}$ may be designed to target $\hat{\pi}_a(x) \propto \pi_a^\beta(x)$ instead, where $\beta \in (0,1)$ is a tempering parameter. The second modification is possible, when $\pi(x) \propto \mathrm{pr}(x)\ell(x)$ and $\pi_a(x) \propto \mathrm{pr}(x)\ell_a(x)$, where $\pi(x)$ is a prior distribution, $\ell(x)$ is a marginal likelihood and $\ell_a(x)$ its approximation. Then, we may use a modified marginal $\hat{\pi}_a(x) \propto \mathrm{pr}(x)(\ell_a(x)+\epsilon)$, with some $\epsilon > 0$. If the likelihood ℓ is bounded, then $\pi(x)/\hat{\pi}_a(x) \propto \ell(x)/(\ell_a(x)+\epsilon)$ is bounded. The latter approach is a specific instance of defensive importance sampling [43], and other generic safe IS schemes may also be useful [cf. 7].

Although we used the global approximation scheme successfully in our experiments, we stress that it may be unstable in some scenarios. This is because the discrepancy between the approximation and the true distribution can increases far from the 'centre' point where the global approximation is constructed. Because the cost of the approximation is independent of m, the overhead from the local approximation scheme is often small in challenging scenarios where m is large.

We used adaptive MCMC in order to construct the marginal chain $(X_k)_{k\geq 1}$ in our experiments. Other adaptation techniques may be applied as well; see for example the review [4]. Note that our theoretical results do not apply directly with adaptive MCMC, unless the adaptation is stopped after suitable burn-in, like we did in our experiments. The results could be extended to hold with continuous adaptation, under certain technical conditions.

An obvious parallelisable alternative to our approach is that, instead of running one Markov chain targeting π_a and then calculating conditional importance sampling type estimates, one may consider running multiple independent pseudo-marginal type MCMC methods in parallel. We note that while valid, this approach may suffer from bias if the MCMC chains are not run long enough, and burn-in of each chain, including possible adaptation, could be seen wasteful. It is well-known that pseudo-marginal methods may be slow mixing when the importance weights have large fluctuations — in the extreme case they may lose desirable rates of convergence such as geometric ergodicity [3, 5]. Large variability of importance weights affects also our method, but the problem is isolated to the correction stage. We believe that MCMC targeting the approximate target π_a may be often easier to design in an efficient manner.

Acknowledgements

The authors have been supported by an Academy of Finland research grants 274740 and 284513. We thank Christophe Andrieu, Anthony Lee, and Arnaud Doucet for many insightful remarks.

Appendix A. Properties of augmented Markov chains

Throughout this section, suppose that K is a Markov transition probability on X and $Q(x,B)$ is a Markov transition probability from X to a space S. We consider here properties of an augmented Markov transition probability $\check{K}$ defined on $\mathsf{X}\times\mathsf{S}$

as follows:

$$\check{K}\big((x,s),\mathrm{d}x'\times\mathrm{d}s'\big) := K(x,\mathrm{d}x')Q(x',\mathrm{d}s').$$

We first state the following basic result.

Lemma 21. *The properties of K and the augmented chain $\check{K}$ are related as follows:*

(i) Let $\mathrm{irr}(K)$ *denote the set of ϕ-irreducibility measures of a Markov transition probability K, then*

$$\varphi_P \in \mathrm{irr}(K) \implies \varphi_{\check{P}}(\mathrm{d}x\times\mathrm{d}s) := \varphi_P(\mathrm{d}x)Q(x,\mathrm{d}s) \in \mathrm{irr}(\check{K}).$$
$$\varphi_{\check{P}} \in \mathrm{irr}(\check{K}) \implies \varphi_P(\mathrm{d}x) := \varphi_{\check{P}}(\mathrm{d}x\times\mathsf{S}) \in \mathrm{irr}(K).$$

(ii) The implications in (i) *hold when* $\mathrm{irr}(K)$ *and* $\mathrm{irr}(\check{K})$ *are replaced with sets of maximal irreducibility measures of K and $\check{K}$, respectively.*

(iii) The invariant probabilities of K and $\check{K}$ satisfy:

$$\mu K = \mu \implies \check{\mu}\check{K} = \check{\mu} \quad \textit{where} \quad \check{\mu}(\mathrm{d}x\times\mathrm{d}y) := \mu(\mathrm{d}x)Q(x,\mathrm{d}y)$$
$$\check{\mu}\check{K} = \check{\mu} \implies \mu K = \mu \quad \textit{where} \quad \mu(\mathrm{d}x) := \check{\mu}(\mathrm{d}x\times\mathsf{S}).$$

The above implications hold also with invariance replaced by reversibility.

(iv) K is Harris recurrent if and only if $\check{K}$ is Harris recurrent.

(v) Suppose $h : \mathsf{X}\times\mathsf{S} \to \mathbb{R}$ is measurable and such that $m_h(x) := \int Q(x,\mathrm{d}s)h(x,s)$ and $(K^n m_h)(x)$ are well-defined. Then, for any $n \geq 1$,

$$(\check{K}^n h)(x,s) = (K^n m_h)(x).$$

Proof. The inheritance of irreducibility measures (i), maximal irreducibility measures (ii), invariant measures (iii), and reversibility are straightforward.

For Harris recurrence (iv), let the probability ϕ_K be a maximal irreducibility measure for K, then $\phi_{\check{K}}(\mathrm{d}x\times\mathrm{d}s) := \phi_K(\mathrm{d}x)Q(x,\mathrm{d}s)$ is the maximal irreducibility measure of $\check{K}$. Let $C \in \mathcal{X}\otimes\mathcal{S}$ with $\phi_{\check{K}}(C) > 0$, and choose $\epsilon > 0$ such that $\phi_K(C(\epsilon)) > 0$, where $C(\epsilon) := \{x \in \mathsf{X} : Q(x,C_x) > \epsilon\}$ with $C_x := \{s \in \mathsf{S} : (x,s) \in C\}$. Notice that

$$\mathbb{P}\bigg(\sum_{k=1}^{\infty}\mathbf{1}\left((X_k,S_k)\in C\right)=\infty\bigg) \geq \mathbb{P}\bigg(\sum_{k=1}^{\infty}\mathbf{1}\left(S_{\tau_k}\in C_{X_{\tau_k}}\right)=\infty\bigg),$$

where τ_k are the hitting times of (X_k) to $C(\epsilon)$. This concludes the proof because $\mathbf{1}\left(S_{\tau_k}\in C_{X_{\tau_k}}\right)$ are independent Bernoulli random variables with success probability at least ϵ. The converse statement is similar.

For (v), it is enough to notice that for any $(x,s) \in \mathsf{X}\times\mathsf{S}$ and $n \geq 1$, it holds that $\check{K}^n\big((x,s),\mathrm{d}x'\times\mathrm{d}s'\big) = K^n(x,\mathrm{d}x')Q(x',\mathrm{d}s')$. □

We next state the following generic results about the asymptotic variance and the central limit theorem of an augmented Markov chain. For $h \in L_0^2(\check{\mu})$, we denote as above the conditional mean $m_h(x) := \int Q(x,\mathrm{d}s)h(x,s)$ and the conditional variance $v_h(x) := \int Q(x,\mathrm{d}s)h^2(x,s) - m_h^2(x)$.

Lemma 22. *Let $h \in L_0^2(\check{\mu})$. The asymptotic variance of an augmented Markov chain satisfies*

$$\mathrm{Var}(h,\check{K}) = \mathrm{Var}(m_h,K) + \mu(v_h),$$

whenever $\mathrm{Var}(m_h, K)$ *is well-defined.*

Proof. Let (X_k, S_k) be a stationary Markov chain with transition probability $\check{K}$.

$$\mathrm{Var}\left(\frac{1}{\sqrt{n}}\sum_{k=1}^{n} h(X_k, S_k)\right) = \check{\mu}(h^2) + \frac{2}{n}\sum_{i=1}^{n-1}\sum_{\ell=1}^{n-i}\mathbb{E}[h(X_0, S_0)h(X_\ell, S_\ell)],$$

by stationarity. For $\ell \geq 1$, Lemma 21 (v) implies

$$\mathbb{E}[h(X_0, S_0)h(X_\ell, S_\ell)] = \mathbb{E}[m_h(X_0)m_h(X_\ell)].$$

We deduce for any $n \geq 1$

$$\mathrm{Var}\left(\frac{1}{\sqrt{n}}\sum_{k=1}^{n} h(X_k, S_k)\right) = \mathrm{Var}\left(\frac{1}{\sqrt{n}}\sum_{k=1}^{n} m_h(X_k)\right) + \mu(v_h),$$

because $\check{\mu}(h^2) - \mu(m_h^2) = \mu(v_h)$. The claim follows by taking limit $n \to \infty$. □

Lemma 23. *Suppose K is Harris-ergodic and $h \in L_0^2(\check{\mu})$. The CLT*

$$\frac{1}{\sqrt{n}}\sum_{k=1}^{n} h(X_k, S_k) \xrightarrow{n\to\infty} N\big(0, \mathrm{Var}(m_h, K) + \mu(v_h)\big) \tag{16}$$

holds for every initial distribution, if one of the following holds:

(i) K is reversible and $\mathrm{Var}(m_h, K) < \infty$.
(ii) $\sum_{n=1}^{\infty} n^{-3/2}\big\{\mu\big(\big[\sum_{i=0}^{n-1} K^i m_h\big]^2\big)\big\}^{1/2} < \infty$.
(iii) There exists $g \in L^2(\mu)$ which solves the Poisson equation $g - Kg = m_h$. In this case $\mathrm{Var}(m_h, K) = \mu(g^2 - (Kg)^2)$.

Proof. The reversible case (i) follows from Lemma 22 and a result of Kipnis and Varadhan [51], which imply (16) from $\check{\mu}$-almost every starting point. Lemma 21 (iv) and and [67, Proposition 17.1.6] guarantee (16) from every starting point.

The case (ii) follows similarly, but relies on a result due to Maxwell and Woodroofe [64], which guarantees (16) from $\check{\mu}$-almost every starting point, if

$$\sum_{n=1}^{\infty} n^{-3/2}\left\{\check{\mu}\left(\left[\sum_{i=0}^{n-1}\check{K}^i h\right]^2\right)\right\}^{1/2} < \infty.$$

Notice that for $n \geq 2$ by Lemma 21 (v),

$$\check{\mu}\left(\left[\sum_{i=0}^{n-1}\check{K}^i h\right]^2\right) = \check{\mu}\left(\left[(h - m_h) + \sum_{i=0}^{n-1} K^i m_h\right]^2\right) = \mu(v_h) + \mu\left(\left[\sum_{i=0}^{n-1} K^i m_h\right]^2\right).$$

Because $(a+b)^{1/2} \leq a^{1/2} + b^{1/2}$ for $a, b \geq 0$ and $\mu(v_h) < \infty$, the claim follows.

For (iii), we first observe that

$$\check{g} - \check{K}\check{g} = h \qquad \text{where} \qquad \check{g}(x, s) := g(x) + h(x, s) - m_h(x) \in L^2(\check{\mu}).$$

Indeed, it is clear that $\check{g} \in L^2(\check{\mu})$ and because $(\check{K}\check{g})(x, s) = (Kg)(x)$,

$$\check{g}(x, s) - (\check{K}\check{g})(x, s) = g(x) - (Kg)(x) + h(x, s) - m_h(x) = h(x, s).$$

The CLT and the expression for the asymptotic variance follow from [67, Theorem 17.4.4]. □

Appendix B. Proofs about CLT and asymptotic variance

Proof of Theorem 8. Whenever $\sum_{i=1}^n W_i > 0$, we may write

$$\sqrt{n}\big[E_n(f) - \pi^*(f)\big] = \frac{n^{-1/2}\sum_{k=1}^n W_k \xi_k(\bar{f})}{n^{-1}\sum_{j=1}^n W_j}.$$

The denominator converges to $c_w > 0$ almost surely, so by Slutsky's lemma it is enough to show that the numerator converges in distribution to $N\big(0, \mathrm{Var}(w_u \bar{f}^*, P) + \pi_a(v)\big)$. This follows from Lemma 23 (i) and (ii), under the conditions (i) and (ii), respectively. □

Proof of Theorem 9. For n large enough such that $\sum_{j=1}^k W_j > 0$ for $j > n$, we may write

$$nv_n = \frac{\frac{1}{n}\sum_{k=1}^n W_k^2 \big(\xi_k(f) - E_n(f)\big)^2}{\big(\frac{1}{n}\sum_{j=1}^m W_j\big)^2}.$$

The denominator converges to c_w^2, and the numerator can be written as

$$\frac{1}{n}\sum_{k=1}^n W_k^2(\xi_k^2(\bar{f}) + D_n^2 + 2\xi_k(\bar{f})D_n), \qquad \text{with} \qquad D_n := \pi^*(f) - E_n(f).$$

The term $n^{-1}\sum_{k=1}^n W_k^2 \xi_k^2(\bar{f}) \to \pi_a(v)$, and because $D_n \to 0$, the remainder terms $D_n^2(n^{-1}\sum_{k=1}^n W_k^2) \to 0$ and $D_n n^{-1}\sum_{k=1}^n W_k^2 \xi_k(\bar{f}) \to 0$, as $|\xi_k(\bar{f})| \le 1 + \xi_k^2(\bar{f})$. □

Appendix C. Proofs about jump chain estimators

The following proposition complements [24, Lemma 1] and [20], which are stated for more specific cases.

Proposition 24. *Suppose (X_k) is a Markov chain with non-degenerate transition probability K and $(\tilde{X}_k)$ its jump chain with corresponding holding times (N_k) (Definition 3). Then, the following hold:*

(i) $(\tilde{X}_k)$ is a Markov chain with transition probability

$$\tilde{K}(x, A) = \frac{K(x, A \setminus \{x\})}{a(x)} \qquad \text{where} \qquad a(x) := K(x, \mathsf{X} \setminus \{x\}) > 0.$$

(ii) The holding times (N_k) are conditionally independent given $(\tilde{X}_k)$, and each N_k has geometric distribution with parameter $a(\tilde{X}_k)$.

(iii) If K admits invariant probability $\mu(\mathrm{d}x)$, then $\tilde{K}$ admits the invariant probability

$$\tilde{\mu}(\mathrm{d}x) := \frac{\mu(\mathrm{d}x)a(x)}{\mu(a)}.$$

In addition, if K is reversible with respect to μ, then $\tilde{K}$ is reversible with respect to $\tilde{\mu}$.

(iv) (X_k) is ψ-irreducible if and only if $(\tilde{X}_k)$ is ψ-irreducible, with the same maximal irreducibility measure.

(v) (X_k) is Harris recurrent if and only if $(\tilde{X}_k)$ is Harris recurrent.

Proof. The expression of the transition probability (i) is due to straightforward conditioning, and (ii) was observed in [24]. The invariance (iii) follows from

$$\begin{aligned}\int \tilde{\mu}(\mathrm{d}x)\tilde{K}(x,A) &= \frac{1}{\mu(a)}\int \mu(\mathrm{d}x)\big[K(x,A) - \mathbf{1}\,(x\in A)\,K(x,\{x\})\big] \\ &= \frac{1}{\mu(a)}\left[\mu(A) - \int_A \mu(\mathrm{d}x)\big(1-a(x)\big)\right] = \tilde{\mu}(A),\end{aligned}$$

and the reversibility is shown in [24]. For (iv) it is sufficient to observe that

$$\forall x\in\mathsf{X}: \sum_{n\geq 1}\mathbb{P}(X_n\in A \mid X_0 = x) > 0 \iff \forall x\in\mathsf{X}: \sum_{n\geq 1}\mathbb{P}(\tilde{X}_n\in A \mid \tilde{X}_0 = x) > 0,$$

which holds because the sets $\{X_k\}_{k\geq 0}$ and $\{\tilde{X}_k\}_{k\geq 0}$ coincide. Similarly, (v) holds because

$$\forall x\in\mathsf{X}: \mathbb{P}(\eta_A = \infty \mid X_0 = x) = 1 \iff \forall x\in\mathsf{X}: \mathbb{P}(\tilde{\eta}_A = \infty \mid \tilde{X}_0 = x) = 1,$$

where $\eta_A := \sum_{k=1}^{\infty}\mathbf{1}\,(X_k\in A)$ and $\tilde{\eta}_A := \sum_{k=1}^{\infty}\mathbf{1}\,(\tilde{X}_k\in A)$. □

We now state results about the asymptotic variance of the jump chain, complementing the reversible case characterisation of Deligiannidis and Lee [20].

Proposition 25. *Suppose $f\in L_0^2(\tilde{\mu})$.*

(i) If K is reversible, then $\mathrm{Var}(f,\tilde{K}) < \infty$ if and only if $\mathrm{Var}(af,K)<\infty$, and

$$\mathrm{Var}(f,\tilde{K}) = \frac{1}{\mu(a)}\left[\mathrm{Var}(af,K) - \mu\big(a(1-a)f^2\big)\right]. \tag{17}$$

(ii) If there exists a function $g\in L^2(\mu)$ which satisfies $g - Kg = af$, then (17) holds, $g-\tilde{K}g = f$ and $g\in L^2(\tilde{\mu})$.

Proof. The reversible case (i) is a restatement of [20, Theorem 1].

Consider then (ii). By Proposition 24 (i), we obtain for any $h:\mathsf{X}\to\mathbb{R}$ with Kh well-defined,

$$(\tilde{K}h)(x) = \frac{(Kh)(x) - \big(1-a(x)\big)h(x)}{a(x)} = \frac{(Kh)(x)-h(x)}{a(x)} + h(x).$$

Consequently, we observe that $g-\tilde{K}g = a^{-1}\big(g-Kg\big) = f$ implying (ii). Because $g\in L^2(\tilde{\mu})$, Lemma 23 (iii) and a straightforward calculation yield

$$\begin{aligned}\mathrm{Var}(f,\tilde{K}) &= \tilde{\mu}\big(g^2 - (\tilde{K}g)^2\big) \\ &= 2\langle g, g-\tilde{K}g\rangle_{\tilde{\mu}} - \langle g-\tilde{K}g, g-\tilde{K}g\rangle_{\tilde{\mu}} \\ &= \mu(a)^{-1}\big[2\langle g, g-Kg\rangle_{\mu} - \mu(af^2)\big],\end{aligned}$$

where $\langle f,g\rangle_{\mu} := \int f(x)g(x)\mu(\mathrm{d}x)$. Similarly, by Lemma 23 (iii)

$$\mathrm{Var}(af,K) = \mu\big(g^2-(Kg)^2\big) = 2\langle g, g-Kg\rangle_{\mu} - \mu(a^2f^2),$$

which allows us to conclude. □

Proof of Theorem 14. Whenever $\sum_{i=1}^n W_i > 0$, we may write

$$\sqrt{n}\big[E_n(f) - \pi^*(f)\big] = \frac{n^{-1/2}\sum_{k=1}^n N_k W_k \xi_k(\bar{f})}{n^{-1}\sum_{j=1}^n W_j}.$$

We shall show below that the CLT holds for the numerator, with asymptotic variance $\sigma_n^2 := \big[\mathrm{Var}(w_u\bar{f}^*, P)+\pi_a(\alpha\tilde{v})\big]/\pi_a(\alpha)$. This implies the claim by Slutsky's lemma, as the denominator converges to $c_w/\pi_a(\alpha)$. For the rest of the proof, let $\tilde{P}$ and $\check{P}$ be the Markov transition probabilities of $(\tilde{X}_k)_{k\geq 1}$ and $(\tilde{X}_k,N_k,W_k,\xi_k(\bar{f}))_{k\geq 1}$, respectively, and let $\tilde{\pi}$ and $\check{\pi}$ be the corresponding invariant probabilities. Note that the function $h(x,n,w,\xi) := nw\xi$ is in $L^2(\check{\pi})$ by assumption (8).

In case (i) holds, also $\tilde{P}$ and $\check{P}$ are reversible by Proposition 24 (iii) and Lemma 21 (iii). Lemma 23 (i) with $K=\tilde{P}$, $\check{K}=\check{P}$, $\mu=\tilde{\pi}$ and $\check{\mu}=\check{\pi}$ implies that a CLT holds for h whenever the asymptotic variance is finite:

$$\mathrm{Var}(h,\check{P}) = \mathrm{Var}\Big(\frac{w_u\bar{f}^*}{\alpha},\tilde{P}\Big) + \frac{\pi_a(\alpha\tilde{v}_{N\xi})}{\pi_a(\alpha)},$$

where, by the variance decomposition formula,

$$\begin{aligned}\tilde{v}_{N\xi}(x) &:= \mathrm{Var}(N_kW_k\xi_k(\bar{f}) \mid \tilde{X}_k = x)\\ &= \tilde{v}(x) + \mathrm{Var}\big(N_k\mathbb{E}[W_k\xi_k(\bar{f}) \mid \tilde{X}_k = x, N_k] \;\big|\; \tilde{X}_k = x\big)\\ &= \tilde{v}(x) + w_u^2(x)\bar{f}^{*2}(x)\big(1-\alpha(x)\big)/\alpha^2(x).\end{aligned}$$

Proposition 25 (i) implies that

$$\mathrm{Var}\Big(\frac{w_u\bar{f}^*}{\alpha},\tilde{P}\Big) = \frac{1}{\pi_a(\alpha)}\left[\mathrm{Var}(w_u\bar{f}^*,P) - \pi_a\Big(\frac{(1-\alpha)w_u^2\bar{f}^{*2}}{\alpha}\Big)\right],$$

which implies $\mathrm{Var}(h,\check{P}) = \sigma_n^2$.

Consider then (ii). Proposition 25 (ii) implies that $g-\tilde{P}g = w_u\bar{f}^*/\alpha$, and $g\in L^2(\tilde{\pi})$. Lemma 23 (iii) implies the CLT, and together with Proposition 25 (ii) lead to $\mathrm{Var}(h,\check{P}) = \sigma_n^2$. □

Appendix D. Proofs about particle filters

Proof of Proposition 17. We follow [1] and prove the statement assuming that the resampling scheme Res satisfies the following property:

$$A^{(1:m)} \sim \mathrm{Res}(\,\cdot \mid \omega^{(1:m)}) \quad \text{satisfies} \quad \mathbb{P}(A^{(j)}=i) = \bar{\omega}^{(i)},$$

for all $i,j\in\{1{:}m\}$. We can do this without loss of generality, because this condition would be satisfied if we applied an independent random permutation of indices to the resampling scheme satisfying Assumption 7. The statement is valid without such randomisation, because both the algorithm and the statements are independent of indexing.

For (i), define the functions $h_T(\alpha_{1:T}) := h(\alpha_{1:T})$, and for $t=T,\ldots,2$

$$h_{t-1}(\alpha_{1:t-1}) := \int h_t(\alpha_{1:t})M_t(\alpha_t \mid \alpha_{1:t-1})G_t(\alpha_{1:t})\mathrm{d}\alpha_t.$$

By Assumption 8, $h_0 := \int M_1(\alpha_1)G_1(\alpha_1)h_1(\alpha_1)\mathrm{d}\alpha_1$ coincides with the desired expectation, and all h_t are (almost everywhere) well-defined if the latter integral is well-defined.

Let us denote $\alpha_{1:t}^{(*)} := \{\alpha_{1:t}^{(i)}\}$ for all $i \in \{1{:}m\}$ and similarly $A_{1:t}^{(*)}$, and observe that for $t = 2, \dots, T$ and $i \in \{1{:}m\}$,

$$
\begin{aligned}
(18) \quad & \mathbb{E}\big[\omega_t^* \bar{\omega}_t^{(i)} h_t(\bar{\alpha}_{t-1}^{(A_{t-1}^{(i)})}, \alpha_t^{(i)}) \mid \alpha_{1:t-1}^{(*)}, A_{1:t-2}^{(*)}\big] \\
& = \mathbb{E}\Big[\mathbb{E}\big[\omega_t^{(i)} h_t(\bar{\alpha}_{t-1}^{(A_{t-1}^{(i)})}, \alpha_t^{(i)}) \mid \alpha_{1:t-1}^{(*)}, A_{1:t-1}^{(*)}\big] \;\Big|\; \alpha_{1:t-1}^{(*)}, A_{1:t-2}^{(*)}\Big] \\
& = \mathbb{E}\Big[\int M_t(\alpha_t \mid \bar{\alpha}_{t-1}^{(A_{t-1}^{(i)})}) G_t(\bar{\alpha}_{t-1}^{(A_{t-1}^{(i)})}, \alpha_t) h_t(\bar{\alpha}_{t-1}^{(A_{t-1}^{(i)})}, \alpha_t) \mathrm{d}\alpha_t \mid \alpha_{1:t-1}^{(*)}, A_{1:t-2}^{(*)}\Big] \\
& = \sum_{j=1}^{m} \bar{\omega}_{t-1}^{(j)} h_{t-1}(\bar{\alpha}_{t-1}^{(j)}).
\end{aligned}
$$

By applying (18) recursively, we obtain

$$
\begin{aligned}
\mathbb{E}[U\xi(h)] &= \mathbb{E}\Big[\Big(\prod_{t=1}^{T-1} \frac{1}{m}\omega_t^*\Big) \frac{1}{m} \sum_{i=1}^{m} \mathbb{E}\Big[\omega_T^* \bar{\omega}_T^{(i)} h_T(\bar{\alpha}_{T-1}^{(A_{T-1}^{(i)})}, \alpha_T^{(i)}) \;\Big|\; \alpha_{1:T-1}^{(*)}, A_{1:T-2}^{(*)}\Big]\Big] \\
&= \mathbb{E}\Big[\Big(\prod_{t=1}^{T-1} \frac{1}{m}\omega_t^*\Big) \sum_{i=1}^{m} \bar{\omega}_{T-1}^{(i)} h_{T-1}(\bar{\alpha}_{T-1}^{(i)})\Big] \\
&= \frac{1}{m} \sum_{i=1}^{m} \mathbb{E}\Big[\omega_1^* \bar{\omega}_1^{(i)} h_1(\alpha_1^{(i)})\Big],
\end{aligned}
$$

which equals to h_0 by a similar calculation as in (18).

The statement (ii) is equivalent with

$$
(19) \qquad \mathbb{E}\Big[U \sum_{i_{1:T} \in \{1:m\}^T} \prod_{t=1}^{T} b_t(i_t \mid i_{t+1}) h(\alpha_{1:T}^{(i_{1:T})})\Big] = h_0.
$$

Similarly as above, we have for $t = T, \dots, 2$ and any $i_{t-1}, i_t \in \{1{:}m\}$,

$$
\begin{aligned}
& \mathbb{E}\big[\omega_t^* \bar{\omega}_t^{(i_t)} b_{t-1}(i_{t-1} \mid i_t) h_t(\alpha_{1:t}^{(i_{1:t})}) \mid \alpha_{1:t-1}^{(*)}, A_{1:t-2}^{(*)}\big] \\
& \quad = \int \sum_{j=1}^{n} \bar{\omega}_{t-1}^{(j)} C_t(\alpha_{t-1}^{(j)}, \alpha_t) \frac{\bar{\omega}_{t-1}^{(i_{t-1})} C_t(\alpha_{t-1}^{(i_{t-1})}, \alpha_t)}{\sum_{\ell=1}^{m} \bar{\omega}_{t-1}^{(\ell)} C_t(\alpha_{t-1}^{(\ell)}, \alpha_t)} h_t(\alpha_{1:t}^{(i_{1:t})}) \mathrm{d}\alpha_t \\
& \quad = \bar{\omega}_{t-1}^{(i_{t-1})} h_{t-1}(\alpha_{1:t-1}^{(i_{1:t-1})}),
\end{aligned}
$$

and application of this recursively leads to (19).

The result (iii) follows also from (19), by first summing over the indices $i_{t+2:T}$. The last result (iv) follows as a special case of (iii). □

References

[1] C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 72(3):269–342, 2010. (with discussion).

[2] C. Andrieu, A. Lee, and M. Vihola. Uniform ergodicity of the iterated conditional SMC and geometric ergodicity of particle Gibbs samplers. *Bernoulli*, to appear.

[3] C. Andrieu and G. O. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. *Ann. Statist.*, 37(2):697–725, 2009.
[4] C. Andrieu and J. Thoms. A tutorial on adaptive MCMC. *Statist. Comput.*, 18(4):343–373, Dec. 2008.
[5] C. Andrieu and M. Vihola. Convergence properties of pseudo-marginal Markov chain Monte Carlo algorithms. *Ann. Appl. Probab.*, 25(2):1030–1077, 2015.
[6] C. Andrieu and M. Vihola. Establishing some order amongst exact approximations of MCMCs. *Ann. Appl. Probab.*, 26(5):2661–2696, 2016.
[7] Y. Z. Art Owen. Safe and effective importance sampling. *J. Amer. Statist. Assoc.*, 95(449):135–143, 2000.
[8] M. Banterle, C. Grazian, A. Lee, and C. P. Robert. Accelerating Metropolis-Hastings algorithms by delayed acceptance. Preprint arXiv:1503.00996v2, 2015.
[9] M. J. Beal. *Variational algorithms for approximate Bayesian inference*. PhD thesis, University College London, 2003.
[10] M. A. Beaumont. Estimation of population growth or decline in genetically monitored populations. *Genetics*, 164:1139–1160, 2003.
[11] A. Beskos, O. Papaspiliopoulos, G. O. Roberts, and P. Fearnhead. Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68(3):333–382, 2006. (with discussion).
[12] S. Bhattacharya. Consistent estimation of the accuracy of importance sampling using regenerative simulation. *Statist. Probab. Lett.*, 78(15):2522–2527, 2008.
[13] O. Cappé, E. Moulines, and T. Rydén. *Inference in Hidden Markov Models*. Springer, 2005.
[14] N. Chopin, P. Jacob, and O. Papaspiliopoulos. SMC^2: A sequential Monte Carlo algorithm with particle Markov chain Monte Carlo updates. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 75(3):397–426, 2013.
[15] N. Chopin and S. S. Singh. On particle Gibbs sampling. *Bernoulli*, 21(3):1855–1883, 2015.
[16] J. A. Christen and C. Fox. Markov chain Monte Carlo using an approximation. *J. Comput. Graph. Statist.*, 14(4), 2005.
[17] P. Del Moral. *Feynman-Kac Formulae*. Springer, 2004.
[18] P. Del Moral, A. Doucet, and A. Jasra. Sequential Monte Carlo samplers. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68(3):411–436, 2006.
[19] G. Deligiannidis, A. Doucet, M. K. Pitt, and R. Kohn. The correlated pseudo-marginal method. Preprint arXiv:1511.04992v3, 2015.
[20] G. Deligiannidis and A. Lee. Which ergodic averages have finite asymptotic variance? Preprint arXiv:1606.08373v1, 2016.
[21] H. Doss. Discussion: Markov chains for exploring posterior distributions. *Ann. Statist.*, 22(4):1728–1734, 1994.
[22] H. Doss. Estimation of large families of Bayes factors from Markov chain output. *Statist. Sinica*, pages 537–560, 2010.
[23] R. Douc, O. Cappé, and E. Moulines. Comparison of resampling schemes for particle filtering. In *Proc. Image and Signal Processing and Analysis, 2005*,

pages 64–69, 2005.
[24] R. Douc and C. P. Robert. A vanilla Rao-Blackwellization of Metropolis-Hastings algorithms. *Ann. Statist.*, 39(1):261–277, 2011.
[25] A. Doucet, N. de Freitas, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, 2001.
[26] A. Doucet, S. Godsill, and C. Andrieu. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statist. Comput.*, 10(3):197–208, 2000.
[27] A. Doucet, M. Pitt, G. Deligiannidis, and R. Kohn. Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika*, 102(2):295–313, 2015.
[28] J. Durbin and S. J. Koopman. Monte Carlo maximum likelihood estimation for non-Gaussian state space models. *Biometrika*, 84(3):669–684, 1997.
[29] J. Durbin and S. J. Koopman. Time series analysis of non-Gaussian observations based on state space models from both classical and Bayesian perspectives. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 62:3–56, 2000.
[30] J. Durbin and S. J. Koopman. A simple and efficient simulation smoother for state space time series analysis. *Biometrika*, 89:603–615, 2002.
[31] J. M. Flegal and G. L. Jones. Batch means and spectral variance estimators in Markov chain Monte Carlo. *Ann. Statist.*, 38(2):1034–1070, 2010.
[32] M. Gerber and N. Chopin. Sequential quasi-Monte Carlo. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 77(3):509–579, 2015. (with discussion).
[33] M. B. Giles. Multilevel Monte Carlo path simulation. *Oper. Res.*, 56(3):607–617, 2008.
[34] P. W. Glynn and D. L. Iglehart. Importance sampling for stochastic simulations. *Management Science*, 35(11):1367–1392, 1989.
[35] P. W. Glynn and W. Whitt. The asymptotic efficiency of simulation estimators. *Oper. Res.*, 40(3):505–520, 1992.
[36] S. J. Godsill, A. Doucet, and M. West. Monte Carlo smoothing for nonlinear time series. *J. Amer. Statist. Assoc.*, 99(465):156–168, 2004.
[37] N. J. Gordon, D. J. Salmond, and A. F. M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings-F*, 140(2):107–113, 1993.
[38] P. J. Green, K. Łatuszyński, M. Pereyra, and C. P. Robert. Bayesian computation: a summary of the current state, and samples backwards and forwards. *Statist. Comput.*, 25(4):835–862, 2015.
[39] P. Guarniero, A. M. Johansen, and A. Lee. The iterated auxiliary particle filter. *J. Amer. Statist. Assoc.*, to appear. Preprint: arXiv:1511.06286.
[40] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, Apr. 1970.
[41] S. Heinrich. Multilevel Monte Carlo methods. In *Large-scale scientific computing*, pages 58–67. Springer, 2001.
[42] J. Helske and M. Vihola. *bssm: Bayesian Inference of State Space Models*, 2016. R package version 0.0.72, available at `https://github.com/helske/bssm`.
[43] T. Hesterberg. Weighted average importance sampling and defensive mixture distributions. *Technometrics*, 37(2):185–194, 1995.

[44] P. E. Jacob, F. Lindsten, and T. B. Schön. Coupling of particle filters. Preprint arXiv:1606.01156v2, 2016.
[45] P. E. Jacob and A. H. Thiery. On nonnegative unbiased estimators. *Ann. Statist.*, 43(2):769–784, 2015.
[46] S. F. Jarner and E. Hansen. Geometric ergodicity of Metropolis algorithms. *Stochastic Process. Appl.*, 85(2):341–361, 2000.
[47] S. F. Jarner and G. O. Roberts. Convergence of heavy-tailed Monte Carlo Markov chain algorithms. *Scand. J. Stat.*, 34(4):781–815, Dec. 2007.
[48] G. L. Jones. On the Markov chain central limit theorem. *Probab. Surv.*, 1:299–320, 2004.
[49] M. I. Jordan. Graphical models. *Statist. Sci.*, pages 140–155, 2004.
[50] G. Karagiannis and C. Andrieu. Annealed importance sampling reversible jump MCMC algorithms. *J. Comput. Graph. Statist.*, 22(3):623–648, 2013.
[51] C. Kipnis and S. S. Varadhan. Central limit theorem for additive functionals of reversible Markov processes and applications to simple exclusions. *Comm. Math. Phys.*, 104(1):1–19, 1986.
[52] G. Kitagawa. Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *J. Comput. Graph. Statist.*, 5(1):1–25, 1996.
[53] S. J. Koopman, N. Shephard, and D. Creal. Testing the assumptions behind importance sampling. *J. Econometrics*, 149(1):2 – 11, 2009.
[54] A. Lee and K. Łatuszynski. Variance bounding and geometric ergodicity of Markov chain Monte Carlo kernels for approximate Bayesian computation. *Biometrika*, 101(3):655–671, 2014.
[55] A. Lee and N. Whiteley. Variance estimation and allocation in the particle filter. Preprint arXiv:1509.00394v2, 2015.
[56] A. Lee, C. Yau, M. B. Giles, A. Doucet, and C. C. Holmes. On the utility of graphics cards to perform massively parallel simulation of advanced Monte Carlo methods. *J. Comput. Graph. Statist.*, 19(4):769–789, 2010.
[57] L. Lin, K. Liu, and J. Sloan. A noisy Monte Carlo algorithm. *Phys. Rev. D*, 61, 2000.
[58] F. Lindsten, R. Douc, and E. Moulines. Uniform ergodicity of the Particle Gibbs sampler. *Scand. J. Stat.*, 42(3):775–797, 2015.
[59] F. Lindsten, M. I. Jordan, and T. B. Schön. Particle Gibbs with ancestor sampling. *J. Mach. Learn. Res.*, 15(1):2145–2184, 2014.
[60] J. S. Liu. Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statist. Comput.*, 6(2):113–119, 1996.
[61] J. S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer-Verlag, New York, 2003.
[62] A.-M. Lyne, M. Girolami, Y. Atchade, H. Strathmann, and D. Simpson. On Russian roulette estimates for Bayesian inference with doubly-intractable likelihoods. *Statist. Sci.*, 30(4):443–467, 2015.
[63] P. Marjoram, J. Molitor, V. Plagnol, and S. Tavaré. Markov chain Monte Carlo without likelihoods. *Proc. Natl. Acad. Sci. USA*, 100(26):15324–15328, 2003.
[64] M. Maxwell and M. Woodroofe. Central limit theorems for additive functionals of Markov chains. *Ann. Probab.*, 28(2):713–724, 2000.

[65] P. McCullagh and J. A. Nelder. *Generalized Linear Models*. Chapman & Hall, London, 2nd edition, 1989.
[66] D. McLeish. A general method for debiasing a Monte Carlo estimator. *Monte Carlo Methods Appl.*, 17(4):301–315, 2011.
[67] S. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Cambridge University Press, second edition, 2009.
[68] T. P. Minka. Expectation propagation for approximate Bayesian inference. In *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, pages 362–369, 2001.
[69] R. M. Neal. Annealed importance sampling. *Statist. Comput.*, 11(2):125–139, 2001.
[70] E. Nummelin. MC's for MCMC'ists. *Int. Statist. Rev.*, 70(2):215–240, 2002.
[71] M. K. Pitt and N. Shephard. Filtering via simulation: Auxiliary particle filters. *J. Amer. Statist. Assoc.*, 94(446):590–599, 1999.
[72] M. K. Pitt, M.-N. Tran, M. Scharth, and R. Kohn. On the existence of moments for high dimensional importance sampling. Preprint arXiv:1307.7975, 2013.
[73] M. Quiroz, M. Villani, and R. Kohn. Exact subsampling MCMC. Preprint arXiv:1603.08232v2, 2016.
[74] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2016.
[75] C.-H. Rhee and P. W. Glynn. Unbiased estimation with square root convergence for SDE models. *Oper. Res.*, 63(5):1026–1043, 2015.
[76] G. O. Roberts and J. S. Rosenthal. Harris recurrence of Metropolis-within-Gibbs and trans-dimensional Markov chains. *Ann. Appl. Probab.*, 16(4):2123–2139, 2006.
[77] G. O. Roberts and R. L. Tweedie. Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. *Biometrika*, 83(1):95–110, 1996.
[78] H. Rue, S. Martino, and N. Chopin. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 71(2):319–392, 2009. (with discussion).
[79] S. Särkkä. *Bayesian filtering and smoothing*. Cambridge University Press, 2013.
[80] N. Shephard and M. K. Pitt. Likelihood analysis of non-Gaussian measurement time series. *Biometrika*, 84(3):653–667, 1997.
[81] C. Sherlock, A. Thiery, and A. Golightly. Efficiency of delayed-acceptance random walk Metropolis algorithms. Preprint arXiv:1506.08155v1, 2015.
[82] C. Sherlock, A. H. Thiery, G. O. Roberts, and J. S. Rosenthal. On the efficiency of pseudo-marginal random walk Metropolis algorithms. *Ann. Statist.*, 43(1):238–275, 2015.
[83] S. S. Singh, F. Lindsten, and E. Moulines. Blocking strategies and stability of particle Gibbs samplers. Preprint arXiv:1509.08362v1, 2015.
[84] A. Tan, H. Doss, and J. P. Hobert. Honest importance sampling with multiple Markov chains. *J. Comput. Graph. Statist.*, 24(3):792–826, 2015.
[85] L. Tierney. Markov chains for exploring posterior distributions. *Ann. Statist.*, 22(4):1701–1728, 1994.

[86] L. Tierney and J. B. Kadane. Accurate approximations for posterior moments and marginal densities. *J. Amer. Statist. Assoc.*, 81(393):82–86, 1986.

[87] M.-N. Tran, M. Scharth, M. K. Pitt, and R. Kohn. Importance sampling squared for Bayesian inference in latent variable models. Preprint arXiv:1309.3339v3, 2014.

[88] M. Vihola. Robust adaptive Metropolis algorithm with coerced acceptance rate. *Statist. Comput.*, 22(5):997–1008, 2012.

[89] M. Vihola. Unbiased estimators and multilevel Monte Carlo. Preprint arXiv:1512.01022v1, 2015.

[90] N. Whiteley. Discussion on Particle Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 72(3):306–307, 2010.

[91] N. Whiteley et al. Stability properties of some particle filters. *Ann. Appl. Probab.*, 23(6):2500–2537, 2013.

[92] N. Whiteley and A. Lee. Twisted particle filters. *Ann. Statist.*, 42(1):115–141, 2014.

[93] N. Whiteley, A. Lee, and K. Heine. On the role of interaction in sequential Monte Carlo algorithms. *Bernoulli*, 22(1):494–529, 2016.

[94] D. J. Wilkinson. Parallel Bayesian computation. In E. J. Kontoghiorghes, editor, *Handbook of Parallel Computing and Statistics*, pages 481–512. Chapman & Hall/CRC, 2005.

Department of Mathematics and Statistics, University of Jyväskylä P.O.Box 35, FI-40014 Univ. of Jyväskylä, Finland

E-mail address, Matti Vihola: matti.vihola@iki.fi

E-mail address, Jouni Helske: jouni.helske@jyu.fi

E-mail address, Jordan Franks: jordan.j.franks@jyu.fi